

**Black and White and Colour**

**Kathleen A. Akins  
Simon Fraser University**

## PART I

### 1.0 Introduction

This paper is the result of a chance remark made by a colleague: “So, you are writing about luminance vision. What exactly *is* luminance vision?” “Really?!” I thought. “How could *anyone* in neuroscience not know *that*?” Now as a child I was always mystified by a common proverb: “Pride goeth before a fall”. If you tried *very* hard and managed to reign in your pride—if pride genuinely ‘*goeth*’—why were you destined to fall? That seemed rather harsh even by the standards of the Old Testament. It was only a few years ago that I suddenly realized that ‘goeth’ does not mean ‘go away’ a discovery that brought absurd relief. Still, it was not until I wrote this paper that I gained a more robust understanding of the proverb. ‘Luminance vision’ is a phrase as common as mud in the vision sciences. But getting a grip on luminance vision (both on what it is and what it is not) is difficult.

The central topic of this paper, as the reader will have guessed, is the nature of luminance vision and the difference between luminance vision and its close cousin, chromatic vision. *Prima facie*, the topic is not very interesting, certainly not to readers whose central interest in the lies in the phenomenology of colour vision. Why should we learn about ‘black and white’ vision when our research concerns colour vision and its phenomenology, topics of greater philosophical interest (objectively speaking)?

The answer is this: On this point, Wilfrid Sellars (Sellars 1956; Sellars 1962) was prescient. When we try to understand the neural processes of visual perception, those that eventuate in our conscious perceptions of the world, we (often) start with an analogy, one taken from our everyday experiences of the physical world. The analogy in this instance is that of ‘black and white versus colour’ a very engrained notion indeed. Since the days of cave paintings, quite literally, we have known that images can be rendered in two ways: In black charcoal, *outline and shade in* your favorite wild beast with a gang of stick men in hot pursuit; now *colour in* the animal with ochre and sienna. Contemporary photography divides up the same way. A

‘black and white’ or monochrome photograph is a rendering of an image “in one colour”, usually along a greyscale, although one could certainly use any other single hue to create a monochrome photograph. Thus old-fashioned sepia-toned photographs are monochrome prints rendered in an orange hue at various levels of darkness. What makes such a photograph useful, however, is that it takes the pattern of *light intensity* in a reflected image—at some level of spatial grain, across some range of illumination, and with a certain degree of contrast—and renders it in a single colour. (Or at least what appears to be a single colour to normal trichromatic viewer. Remember that a black and white printer can use black ink or a triplet of ink colours. A tri-colour printer produces ‘monochrome’ prints as long as the trichromatic viewer is unable to see any colour differences or specific colours within the printed image.) In a monochrome photograph, we see what is portrayed in virtue of a rendering of image intensity. A *colour* photograph, on the other hand, also renders light intensity but to this is added *hue*, wavelength contrast within the image that *is* accessible to the human visual system. Prior to colour film and printing, for example, black and white photographs were often hand-tinted by painting over them with translucent coloured pigments. Digital photography programs have similar functions. You can ‘paint’ with virtual translucent colour over the black and white image using a virtual brush. So in all cases, from cave paintings to digital photo advertisements, colouring a black and white image adds, well, *colour*—from which one can infer the colours of objects in the world. This is how the black and white or colour distinction became so embedded in the cognitive psychology of most citizens of the contemporary world.

Like Sellars, I suspect that the distinction between ‘black and white’ and ‘colour’ was applied, in the first instance, to physical media—illustrations, photographs and what not—and that the same terminology was then borrowed to describe visual sensations, to one aspect of our visual experience when we inspect black and white or colour images or when we find our way around at night. This view is highly controversial, of course, but fortunately it is not a point on which much hangs at the moment (I hope). What *is* clear, however, is this. If we think of the neurophysiological distinction between luminance and chromatic systems of vision as one of ‘black and white’ and ‘colour’, this *is* an analogy. No one thinks that there is literally a black and white image that is passed from retina to cortex with a quick stop at the LGN. There is no parallel system that conveys the coloured pigment, spatially arranged, to be paired with the

monochrome image wherever it ends up. This is an analogy, one meant more or less literally depending upon who uses it—and, given the particulars of its usage, it will turn out to be more or less apt. It is also clear that the analogy is deeply entrenched. It is very difficult to imagine the workings of the visual brain along any other lines except the division between the ‘black and white’ and the ‘colour’ of public images. When one first learns that the ganglion cells in the retina are of two types, ‘chromatic’ or ‘luminance’ cells, it is natural to think that here too ‘black and white’ and ‘colour’ is the essence of the divide: luminance cells encode light intensity (i.e. brightness or darkness) and chromatic cells encode, well, the *other* dimension of light, wavelength or hue. If not that, what would the nature of the division be?

The central task of this essay is to pry the reader (not to mention, the author) out of the analogy’s firm grip. To do so, we will look at the case of the rod achromat, a person who has only one type of photoreceptor, the rods, and whose visual experience depends upon luminance information alone. Although the normal trichromat also has a rod-based visual system—night vision—we will be looking at the pure case of luminance vision, a person who has, and has always had, only luminance vision. By the end of the essay, it should be clear why this was a good place to begin: The rod achromat’s experience is quite different from what most people would imagine it to be. Hence, our own experience of night vision may not be quite what we imagine either.

As the reader will have guessed, the nature of luminance vision is only the *prima facie* topic of this paper. The hidden agenda is a philosophical one, about the nature of visual experience. The picture of luminance and chromatic processing that emerges, with a restructuring of the ‘black and white or colour’ divide, is of two systems that function in analogous and complementary ways to discern the multiple features of the visual world. The claim is not that luminance and chromatic systems work in exactly the same ways, i.e. they instantiate the same algorithms—but that the two systems use comparable (and often common) mechanisms to perform the multitude of visual functions that comprise human vision. They are intertwined systems, both of which are concerned with the broad goal of *seeing the distal world*. The point of dismantling the analogy, then, is *to make room for chromatic processing*. If a sharp black and white photograph shows you more or less what you would see if you looked at the

same scene in person—if *that* is what you get from luminance processing—then there is only one thing left for chromatic processing to contribute: The colours. In other words, this common analogy, between black and white or colour, is a hindrance to understanding the natural fault lines of human visual processing. And we cannot understand the phenomenology of vision if we do not have these fault lines firmly in place.

### **1.1. Luminance Vision in the Rod Achromat**

What is it like to be a human rod achromat?

A rod achromat is a person who lacks all three of the cones in a normal human trichromatic retina (a ‘complete’ achromat) or a person who lacks these cones functionally if not anatomically (i.e. the retinæ of some ‘achromats’ contain cones but they do not contribute to normal vision.) In short, the rod achromat lacks the human system for ‘daylight’ vision, trichromatic vision.

What the rod achromat retains, however, is a virtually normal rod system for low light or night vision. In the human trichromatic retina, there are two main pathways for luminance information, one from the rods and one from the cones. (This often comes as a surprise to people who have been taught since grade school that “rods are for ‘black-and-white’ and cones are for ‘colour’”).) Both luminance systems use the same outgoing pathway from the retina, the magnocellular pathway. But because the rods and cones function under different levels of illumination, their use constitutes a sort of ‘timeshare’ arrangement, depending upon the light level: In daylight, the cones send luminance signals to cortex via the magnocellular pathway, while at night, the magnocellular pathway carries luminance information from the rods. (It’s a bit like students who ‘double bunk’ for lack of money: the ‘day crew’ studies by day and sleeps at night; the ‘night crew’ sleeps during the day and occupies the desks at night.) In the retina of the rod achromat, the rods still use the magnocellular pathway despite the fact that there are no cones present to take over the magnocellular pathway in bright lighting conditions. So, surprisingly, given all the different ways that achromatic rod vision could have been organized, an achromat’s retina has roughly the same arrangement as our own minus the cones.

The central visual problem for rod achromats is that the rod system does not function well under daylight conditions. For one, the rod system ‘saturates’ under normal daylight conditions. Each rod absorbs so many photons in bright light that the photoreceptors are bleached of all pigment — and without adequate pigment, the rods no longer respond to light. For the rod achromat, then, sight under daylight conditions is very much like the experience the average human trichromat has when someone suddenly flicks on the bedroom light in the middle of the night. Certainly this *hurts* but it also renders the newly awakened subject entirely blind. This is more or less the constant state of the rod achromat in bright sunlight. The rod achromat’s photosensitivity explains why they prefer darkened rooms and deep shadow, and why even under low light conditions rod achromats wear sunglasses.

Second, the rod system is a highly convergent system: it pools together the signals from many different rods in order to maximize photon catch. To see anything at all at night, one needs a system that makes the best possible use of the miniscule amount of available light. However, this pooling of rod signals also decreases the spatial resolution of the system, i.e. the ability to distinguish two distinct but nearby points. (The higher the visual acuity or spatial resolution of a system, the closer together two points can be and still be seen as being distinct.) So, the rod system, in both achromats and trichomats, has far lower spatial resolution than the cone or daylight system of the trichromat. Just as a myopic trichromat is aided by large print or magnified illustrations, so too is the rod achromat.

The difference in spatial acuity between the trichromat and the achromat is actually a bit more complicated than this. As all mothers know, visual acuity is relative to the ambient light level. This is why, when as a child you sat in the dark reading, your mother probably said (sweetly) “TURN ON THAT LIGHT!” Spatial acuity gets better with increased illumination and this is true for both the trichromat and the rod achromat. The difference between these two systems is the absolute light levels at which rods and cones saturate. For the trichromat, during the day, visual acuity increases with light level and suddenly drops off when the photoreceptors saturate in intensely bright light. The same holds true of the rod system but saturation is reached at much lower light levels. The net effect is that the spatial acuity of rod vision is maximized

under ‘mesopic’ conditions, during dawn and dusk or under the illumination of a full moon at night. Under mesopic conditions the rod achromat has the greatest spatial acuity, indeed the same visual acuity as the trichromat in similar circumstances. During daylight hours, however, the rod achromat is ‘all but blind’. But this is not a result of the spatial acuity of the rod system.

So what is it like to be a rod achromat? Dr. Knut Nordby was both a rod achromat and one of the vision scientists responsible for understanding the physiology of rod achromacy (Skottun, Nordby et al. 1982; Hess and Nordby 1986; Greenlee, Magnussen et al. 1988). He described his visual experience as follows:

Trying to explain to someone with normal, or nearly normal, colour-vision what it is like to be totally colour-blind, is probably a bit like trying to describe to a normally hearing person what it is like to be completely tone-deaf, i.e. not possessing the ability to perceive tonal pitch and music. My task, though, is probably a bit simpler than the case of the tone-deaf, since practically everyone has had experiences of achromatic (i.e. colour-less or black & white) or monochrome pictures and renderings, and certainly must have witnessed the gradual disappearance of colours when darkness sets in.

A first approximation, then, in explaining what my colour-less world is like, is to compare it to the visual experiences people with normal colour-vision have when viewing a black & white film in a cinema or when looking at good black & white photographic prints (good here meaning sharply focused, high contrast with a long grey-scale, as in crisp, high quality, glossy, technical prints).

This, however, is only part of the story because I have so far only dealt with the achromatic aspect of my perception. To get a fuller understanding of my visual world one must, in addition to my colour blindness, also take into account my light aversion (i.e. hyper-sensitivity to light) and my reduced visual acuity. (Nordby 1996)

This description sums up Knut Nordby’s view of his own visual experience, one that accords well with commonly made inferences about achromatic experience from the third person point of view. First, given that both a trichromatic retina and the achromat’s retina have very similar rod systems, what the achromat sees at night is probably very close to what the trichromat sees at night: in both cases, the subject sees the world via input from the rods alone. Second, the night vision of a trichromat is commonly described as seeing ‘in black and white’ or like looking at a monochrome photograph. Making allowance for the lack of illumination at night and the relative ‘fuzziness’ of rod-based vision even under optimal (mesopic) lighting conditions, what trichromats see at night is very much what we see when looking at a

monochrome photograph or watching an old black and white movie. By transitivity, then, the rod achromat sees just what you, a slightly myopic trichromat, would see were you to watch an old black and white film without your glasses but while experiencing a photosensitive headache. And speaking for myself, this is something I *can* easily imagine. (Surprise! A rod achromat is not an alien-life form. Who would have guessed?)

In the next section, I will begin to explain why neither of the above two inferential steps are good ones. The trichromat's experience of the visual world at night is not like the rod achromat's experience of the world during the day; nor is trichromatic night vision like looking at a black and white movie or photograph.

## 1.2 Luminance Information

Like many terms commonly used in science (e.g. 'electron' 'energy' 'power'), we often use 'luminance' without remembering (if we ever knew) its explicit definitions. 'Luminance', we all know, has something to do with 'the amount of light'. But if you look up the definition of 'luminance', you will find the following puzzling statement: luminance is the radiant intensity of light as filtered by the human photopic luminance function. In turn, if you look up 'photopic luminance function', you will learn that this is a model (for the normal trichromatic human observer, adapted to photopic conditions) of the probability that an individual photon will be absorbed, expressed as a function of the wavelength of the stimulus. Soooo....after suitable rumination, it seems that 'luminance' refers to the amount of light, at each wavelength, your visual system will absorb. Uh huh. While technically correct, there is something unsatisfying about this definition. We want to know what a luminance system *does*. But this definition does not explain the *function* of a luminance system or distinguish it from the obvious alternative, the chromatic system. Nor does it tell us what luminance vision represents. It appears to say only that a luminance system absorbs photons—presumably, for the purpose of seeing although even this is not a part of the definition. However unsatisfying this definition may seem, I've come to realize that it has profound consequences.

Consider first that all photopigments on earth function in fundamentally the same way.



Importantly, each photopigment responds to light across a restricted range of wavelengths, what is commonly known as the receptor's spectral range. Within this range, the response of each photopigment is wavelength sensitive: a light stimulus of the same intensity will be absorbed with a greater or lesser probability depending upon the wavelength of the stimulus. Fig. 1 illustrates the response curves of the three human cones. Relative to a fixed intensity or amplitude of light—and that is the part to always remember—the graph illustrates the probability that a photon will be absorbed at each wavelength across its spectral range. At the apex of the curve is photopigment's 'preferred' wavelength, the wavelength of light that will result in the greatest light absorption. For example, peak light absorption for the human long wavelength (or L) cone occurs at 470 nm. Importantly, though, this is only the receptor's preferred wavelength; it responds across the entire spectral range although to lesser extent. What makes one photoreceptor different from another is the photopigment it contains, and this in turn means a difference in the spectral range over which each photoreceptor responds.

What differentiates cones from rods is primarily their sensitivity to light, how much light is required to affect a response. Again, if one thinks in the old terms of 'rods are for black and white' and 'cones are for colour', this might come as a surprise. Above, in the case of the achromat, we saw that cones are responsible for luminance signals under photopic conditions. But does it follow that rods are equally capable of producing colour vision? Yes and no, replied the philosopher. Rods, just like cones, respond within a specific spectral window; rods are wavelength sensitive in exactly the same way as cones. The primary difference between rods and cones, as I have said, is the energy required for photon absorption: rods require far less energy and are thus ideal for low light conditions. However, despite their greater sensitivity, the absolute photon catch of rods is still markedly lower than that of cones. This is why rod systems are convergent: they must pool the signals of multiple rods in order to achieve a good signal-to-noise ratio. If yet another type of rod were added uniformly throughout the retina—and one must have at least two types of receptors to discriminate wavelength—this would *halve* again the already poor spatial resolution of night vision. In the dark of night, it is thus the low photon catch of the rods that disqualifies rods for participation in colour vision<sup>1</sup>.

---

<sup>1</sup> Note that this does not rule out have a specialized area of the retina in which rod input is used for chromatic processing. This would be used much like the 'bucket' function in painting

Still, there is nothing in the function of rods that intrinsically precludes them from chromatic processing and the question of whether (and what) rods might contribute to chromatic vision has been an active one since the 1960's (Stabell and Stabell 1965; Stabell 1967; Stabell 1967). It used be thought that rod and cone systems were functionally segregated by light level, i.e. the cones 'shut off' at precisely the level of illumination at which the rods 'come on' and vice versa. We now realize that this is false. At dawn, dusk and under the light of a full moon (under *mesopic* conditions), the mammalian visual system contains a sub-population of rods that contribute to chromatic vision (Buck, Thomas et al. 2006; Cao, Pokorny et al. 2008; Li, Chen et al. 2010; Pang, Gao et al. 2010; Cao, Pokorny et al. 2011). So the world continues to look coloured even when the photon absorption of the cones is compromised. Moreover, very recent experiments suggest that even under scotopic conditions (in the dark of night without starlight), rods feed into the S cone chromatic pathway—which explains why the predominant colour of night, for the trichromat, seems to be blue (Field, Greschner et al. 2009). In dim lighting, we *do* see colours partially as a result of *rod* processing. In retrospect, given the similarities between rods and cones, it is not surprising that rods and cones work together under mesopic and scotopic conditions to encode colour. But this co-operative function is predictable only against a general understanding of photoreceptor function.

So human vision has (at least) two major luminance systems, a photopic (bright light) system that depends upon cone input and a scotopic (low light) system that sums rod signals. *Importantly, neither of these luminance systems—indeed no biological luminance system—encodes light intensity per se.* This is in stark contrast (sorry) to a black and white photograph (at least one printed from colour-corrected black and white film) in which the intensity value of light at each point in the photographic image is represented using the greyscale. (One source of confusion for the reader may be that monochrome images, which represent light intensity, are sometimes called 'luminance images' in computer science and artificial intelligence circles.)

This point is the flip side of one familiar to all researchers of colour vision: A visual

---

programs in which colour is added to an object, not pixel by pixel, but given the outlines of the object to be painted.

system with a single receptor cannot discriminate between two stimuli that differ only in wavelength. In the above graph of cone function, recall that the graph plots the probability that a photon will be absorbed against the wavelength of light relative to *a set intensity of light*. Alter the intensity of the light stimulus and the probability that a photon will be absorbed (at a specific wavelength) is altered as well. So the photon absorption of any receptor is a function of both wavelength and intensity. Receptor response does not indicate or provide information about either property independently of the other. As I said above, for colour researchers, this is a well-known—and one might even say ‘shopworn’—fact: A visual system with a single receptor (or, what comes to the same thing, without the ability to compare different photoreceptor signals) is ‘colour blind’. But what is sauce for the goose is sauce for the gander. The same moral holds, *mutatis mutandis*, for the *intensity* of the light stimulus. If a single photoreceptor conflates the wavelength and intensity of the light stimulus, then each photoreceptor is *intensity blind* as well. This is a fact we don’t hear repeated nearly as often. Without a signal comparison between two different types of photoreceptors, intensity cannot be distinguished from wavelength. Thus, a luminance system, which has only one kind of receptor, is just that: intensity blind.

In 2010, there was an art exhibition in Berlin by the design firm Carnovsky that provided a brilliant demonstration of this fact—and of the nature of luminance vision in general. For this exhibit, entitled ‘RGB’, Carnovsky produced several different wallpapers covered in 19th century illustrations of various species (some of which are even recognizable as the species which they represent). Each animal is rendered by a line drawing—this is important as we’ll see—and printed in one of three colours from the standard printer’s RGB palette of cyan, yellow and magenta. Although the wallpaper is printed digitally, it looks like a screen print with three colours of creatures layered one upon the other. Under natural illumination or any light source that approximates a uniform spectral power distribution (‘white’ light), the coloured figures are clearly visible to the human trichromat. The exceptions are the yellow figures on the wallpaper which can be quite difficult to see especially when overlaid with other creatures, a common problem with yellow figures. (Even though the Carnovsky figures are rendered with a yellow pigment that is just as bright as the other colours, perceptual yellow can never be made as bright as the other colours in human vision.) While the wallpaper is pleasant enough in daylight, the interest of the exhibit really lies in what happens when the wallpaper is illuminated by one of

three coloured lights (Figures 2b, 3a and 3b). When a filtered light is switched on, the entire room is suffused with colour and the illustrations themselves now appear as monochrome images rendered in red, green or blue —or what one might call ‘black and red’, ‘green and black’ or ‘blue and black’ (as opposed to black and blue). Some creatures simply disappear, while those that remain appear in very dark shades of the illuminant colours, almost black. Switching between the coloured lights produces dramatic differences in the visibility of the various creatures. For example, under the red light, the blue creatures are visible but the magenta and yellow ones disappear. Under the blue light almost all of the animals are visible but the yellow creatures, previously invisible, now pop out (as black!); the other creatures appear as more misty grey background figures.

Let’s take a close look at what is going on in this exhibit. (In figuring out this exhibit, I found it helpful to pick out three figures, one in each colour of ink, from the original wallpaper and then to compare their appearance under each of the three coloured lights. So let the fox be our magenta figure, the alligator be cyan or blue, and the large cockroach be yellow. I know. What cockroach? But it is there, overlaid upon the elephant, visible only under the blue light). There are two central ‘effects’ that create the magic of the RGB exhibit. First, the display uses *spectral filtering* to its best advantage, a ‘trick’ that every natural system of vision ‘learns’ to employ over the course of evolution. In the case of the rod achromat, this ‘filter’ is on the receiver end of things: with only one photoreceptor, the rods, visible light is limited to the spectral range of that single receptor. In the RGB exhibit, the normal trichromat observes a room illuminated by one highly filtered light source, by the red, green, or blue light. Here, visible light is limited to an artificially small window by the filtered light source. That is, for us as trichromats, visible or ‘effective’ light ranges from 370 to 660 nm, a spectral range of roughly 300 nm. But under the filtered lights of the exhibit, all light within the room is restricted to a narrow band of light about 60 nm in width. For the trichromatic viewer, then, spectral bandwidth is restricted by the sender not the receiver. Under the red, blue or green lights, whatever the trichromat sees is made visible by a single narrow spectral band of light, be it red, blue or green, reflected from the surfaces within the room. In effect, then, trichromatic observers have reduced spectral range very much like the restricted range of the rod achromat. In fact, one can think of the three lights as producing (very roughly) *functional monochromats*, each with

only a short (blue), medium (green) and long (red) cone/photoreceptor.

From this perspective, what we'll call the perspective of the 'Carnovsky monochromat', it is clear that the strength of the returning signal is strongly *colour dependent*. Each of the three inks (red, blue or yellow) absorbs and reflects light continuously across the normal spectrum of visible light; their surface spectral reflectance or SSR is a continuous function. Yet each ink absorbs and reflects light in a wavelength selective manner. Blue objects appear blue because they reflect *more* short-wavelength light given a light source that emits light at a uniform intensity across the wavelength spectrum. Of course, if the light source does not emit 'blue' light, then there is no light for a blue object to reflect. In such a case, say given a predominantly 'red' light source, a blue object will appear black. This is one of the central principles used in RGB exhibit. Each ink is colour selective: there is a certain range of light wavelengths that it reflects preferentially. Hence the colour of the light can be chosen so as to maximize or minimize the visibility of each colour of ink and, by extension, the visibility of the creatures rendered in each ink colour.

Take, for example the blue alligator. Even though the alligator's cyan pigment absorbs light continuously across the visible spectrum light, cyan pigment *reflects* far more blue light than it absorbs (which is why it appears blue). Conversely, it *absorbs* far more red light than it reflects (why it does *not* appear red). The blue alligator, illuminated with red light, will thus appear black for it reflects almost no red light (and there is no other light to reflect). In this exhibit, blue light produces the most interesting effects. Neither the yellow nor the red figures will reflect much blue light, so both will 'pop out' under the blue light, a particularly good effect for the all-but-invisible yellow creatures (in daylight). But the yellow figures will also reflect a bit more blue light than the red ones. For this reason the red figures will appear both darker and closer, while the yellow ones will appear as more hazy background figures. The cockroach is the exception that proves the rule: it appears *in front of* the elephant under blue light. This is because darkness is also function of the level of detail in the line drawings and thus how much pigment is used. The more yellow pigment per unit of area, the closer the lines of drawing, the darker the yellow creature will appear. Here the cockroach is much more finely rendered than the elephant, hence the cockroach appears to be in front, the darker elephant behind.

More succinctly, the RGB exhibit uses narrow bandwidth filters to re-create the monochromat's world, a world in which perceptual 'lightness' is a function of *both* intensity and the predominant wavelength of the reflected image. With only one photopigment, rhodopsin, the rod achromat's visual world varies along a single visual dimension. So too do the perceptions of the 'functional monochromat' who views the Carnovsky world of illustrated figures under coloured light. Still, there is a crucial difference between a trichromat who views a Carnovsky exhibit under filtered light and a monochromat who views the natural world under sunlight. There is no escaping the fact that, for the trichromat, the RGB exhibit appears in shades of red (or green or blue.) The trichromat sees the light and the wall *as coloured*, as having a particular hue, even if the light and every surface are monochrome, i.e. even though they have the *same* hue. This is not information that the rod monochromat could possibly have, that the illuminant has a particular predominant wavelength as does the light reflected from every surface. We must assume therefore that monochromat's experience is not 'coloured' red *or* blue *or* green and that, in all likelihood, it differs from our experience in this crucial way. This brings us to the second reason why the RGB exhibit works so well.

The second reason why the RGB exhibit is so effective is that the combination of each coloured light and the three ink pigments are designed to enhance (or diminish) *luminance contrast*. When the wallpaper is bathed in red light, for example, only red light can be reflected back from the illustrations. Similarly the white wall, which normally reflects light equally across the entire spectrum, reflects only red light. So, to the trichromat, the wall appears red. As we said, for the trichromat, under a red illuminant, every thing that is visible appears in shades of red from bright red to red-black. But what is *visible* against a bright red wall? A magenta figure (e.g. the fox) will reflect a large percentage of red light. But it will not appear as a red fox against a red wall. A red fox on a red wall lacks sufficient contrast to be seen at all. The same holds true for all of the magenta figures. Paradoxically, under the red illuminant, figures rendered in the *blue* ink will be the most visible. A blue figure reflects very little red light under any lighting conditions, hence it will now reflect very little light *at all*. The blue alligator thus appears as a *black* figure against a red wall. Finally, the yellow figures will now be entirely invisible. We are not told the spectral power distributions (SPD) of the coloured lights used in the exhibit. But

suppose that the red light source contained some ‘yellow’ light and that the yellow pigment reflects a bit of red light in addition to yellow light. This lack of visual contrast would render the yellow figures invisible.

The two principles used by the Carnovsky exhibit—of spectral filtering and luminance contrast—mirror two of the most important principles of vision. In fact, this is why the RGB exhibit works so well on us. First, from above, every known photopigment acts like a wavelength filter, responding to light as a function of both wavelength and intensity. Two different pigments may produce profoundly different levels of excitation in response to one and the same reflected figure. In the evolution of any visual system, the type of photopigments/filters in place will have had a direct effect on visibility within the environment and hence on the species ability to *see* its predators, find sustenance, determine the fitness of mates and so on. (In fact it is hard to imagine many physiological facts that would play as important a role as photopigment sensitivity in the general fitness of a species.)

Figure 4 (Gegenfurtner 2003) demonstrates the human case, the outcome for the majority of our species. Each illustration shows how the S, M or L cones filter a natural image, the photon catch for each of the three cones given the same reflected image. The original colour photograph depicts a group of fruits and vegetables (Fig. 4(a)), most of which reflect light predominantly from the middle- to long-portion of the spectrum, the yellows and oranges. (Blue vegetables are rare, not to mention somewhat unsettling.) Note, for example, the banana and orange bell pepper in the original colour photograph. Now compare their luminance images as filtered by the M and S cones (Figures 4(e) and 4(d) respectively). The M cone is preferentially sensitive to middle wavelengths with its peak preference in the yellow range. So the brightest objects in Figure 4(e), which illustrate net photon absorption by the M cones, are the banana and a lemon. In Figure 4(d) shows the intensity image as filtered by the S cones: here the orange pepper is black and the banana is dark gray. This is because S cones are highly sensitive to ‘blue’ light and insensitive to the longer ‘red’ wavelengths. So in the S cone luminance image, the redder the fruit or vegetable, the darker it will appear. In sum, the amount of ‘effective’ light reflected from an object, the intensity of light as filtered by one photoreceptor or another, depends upon the *spectral* facts of the environment—the colour of both the object and of the light source—and on

the spectral sensitivity of the photoreceptor at issue. This is why the definition of luminance provided at the outset ('luminance is the radiant intensity of light as filtered by the human photopic luminance function') while seemingly empty carries such weight. The information available to any visual system is as much a function of the wavelength of light reflected from distal objects, as it is a function of its intensity. For a luminance system, object colour matters just as much as object lightness or darkness.

Second, the primary concern of evolution in vision—i.e. what natural selection hinges upon—is *what the organism can see*, the *visibility* of relevant objects, not which objects reflect the greatest or least amount of light. Whatever else, the viewer must segregate an item of interest from its background. So at bottom luminance vision requires the registration of luminance *contrast* between an object and its background. It does not matter whether, for this particular visual system, the object has positive or negative contrast with its background—or whether the luminance contrast arises as a function of genuine intensity differences between the object and its background or because, while the figure and ground reflect the same intensity of light, the spectral sensitivity of the cones 'creates' luminance contrast given their difference in colour. 'It's *all* good!' as they say, as long as we are able to distinguish between an object and its background, as long as there is contrast.

A more recent Carnovsky exhibit nicely illustrates this principle of 'visibility by contrast'. In their second large installation, the space to be 'papered' contained two mirror-image rooms, a design feature that the Carnovsky designers wanted to exploit. Both rooms were papered with a jungle scene in which its inhabitants were obscured by dense foliage when viewed under an even SPD illuminant. One room, designated the 'positive' room, was papered with coloured figures on a white ground just as in the RGB exhibit. The second 'negative' room portrayed coloured figures on a black background when viewed under normal lighting conditions. Looking at one part of the display, say the hidden elephant in Fig. 5 (a), one further difference is apparent: the 'positive' room has a blue elephant, hence a blue-on-white image, and the 'negative' room (Fig. 5(b)) has a red elephant hence a red-on-black image. Once the red light is turned on, the rooms demonstrate their intended yin/yang nature. The first room reveals a *black* elephant on a bright *red* wall (a positive image) while the second room shows a bright *red*



elephant on a *black* wall (a reverse image of the other room). Of course, in some sense, no one should be surprised to see a red elephant on black in the negative room. After all, it *is* an illustration of a red elephant drawn on a black background. The surprise is the positive room: the blue elephant on a white wall appears to be a *black* elephant against a *red* background. Added to this ‘reversal’ is a very nice feature of both rooms. In both the positive and negative configurations, the elephants now stand in plain sight without one bit of flora to hide them, a very nice illusion. Note that both of the rooms, under red illumination, produce monochromatic visual images, one the reverse polarity of the other, both of which allow us to see the elephant without difficulty despite the reversal of contrast. Again, it’s luminance contrast that makes visible the figures, not absolute luminance.

Before ending this section, let me explain in slightly different terms, *what* the Carnovsky monochromat loses under the narrow-band illumination. Above, I explained the disappearance of certain figures (in the wallpaper) under coloured lights as the result of decreased wavelength contrast. Obviously, if a red figure is against a red background, one cannot see it. But it is important to realize that a Carnovsky monochromat also loses *intensity information*, what we think of as the ‘black and white’ of the original wallpaper. Compare three images derived from the Carnovsky exhibit (Fig. 2: (a) A *full colour* photograph of the wallpaper under daylight (Fig. 2a.) (ii) An *intensity image* of that same wallpaper, i.e. an image of the wallpaper rendered in greyscale (Fig. 7a); and (iii) A *luminance image* of the wallpaper or an photograph of the wallpaper illuminated by the *green* light source which has been rendered as a monochrome black and white image (Fig. 7b). In the original colour image, all three ink pigments—cyan, magenta, and yellow—reflect roughly the same overall amount of light. So under daylight, all of the figures have roughly the same intensity, i.e. appear equally dark to the trichromatic viewer in the greyscale intensity image. (Some of the creatures *do* look darker than others because they are rendered using many more lines. See the paragraph below for an explanation.) In the luminance image, however, the figures have a range of luminance values, from very dark to shadowy to very faint images. There are also figures that have disappeared, which no longer have any luminance value at all. You can see this clearly if look just to the left of where the walls meet in the luminance and intensity images. In the *intensity* image, there are so many figures rendered one on top of the other that it is difficult to extricate, visually, any one of them. But in the same

area of the *luminance* image, there is one predominant figure: a giant squid. You can also clearly see several surrounding figures: a conch shell, a sea star, the lion's tail and, in shadow, a large fish with a spikey dorsal fin. By looking at the *coloured* wallpaper, you can also see what creatures have been 'disappeared'—i.e. a huge coiled snake and, by the floor, a large mammal akin to a walrus. Above, I explained why certain figures disappear in terms of the *colour* contrast: A red figure against a red wall is invisible. But this point can also be put in terms of light intensity: *there is less luminance (in the luminance image) than intensity contrast (in the intensity image)*. Any narrow-band spectral filter will eliminate all but a small range of wavelengths in the image. But if a spectral filter, like the rods, leads to a loss of light, it leads to a loss of *intensity information as well*.<sup>2</sup> This is why a rod achromat also loses 'black and white', the intensity information of the retinal image.

It is worth emphasizing in the present context how difficult it is 'simply imagine' the achromat's point of view—the effect of a spectral filter on a natural image and the probable consequences for achromatic visual experience. The results are too complex *and* unintuitive. Even people who work with colour and light for a living, the graphic designers at Carnovsky, could not have imagined exactly what would happen when, with their wallpaper newly affixed to the wall, they turned on the first of the coloured lights. In fact, the first RGB images required extensive experimentation with different ink pigments and filtered lights even though the scene was highly artificial, a series of uniformly pigmented line drawings against a uniformly white background (*personal correspondence with Francesco Rugi of Carnovsky Designs*). This same failure of imagination is no less likely for vision scientists or (dare I suggest) philosophers of vision science. A seasoned psychophysicist of colour vision would not be able to predict, accurately and in detail, the appearance of an arbitrary natural scene relative to a specified filter. This is why Gegenfurtner (2003) was allowed to include three images of a fruit and vegetable arrangement in an article in *Nature Reviews Neuroscience*. It is one thing to understand the theoretical principles of spectral filtering and 'visibility by contrast', yet another to imagine the

---

<sup>2</sup> I do not mean to suggest that the luminance image in Figure is an accurate depiction of the rod achromat's situation — either the luminance information available to the rod achromat much less what the rod achromat would see when looking at the Carnovsky exhibit under natural daylight.

concrete, particular results of their application to a natural scene. So a realistic exploration of achromatic vision would begin with detailed image analysis: To see what the achromat can see, that is, we would start with a series of natural images, apply the rod luminance function to each image pixel-by-pixel, and then do a statistical analysis of the resulting set of images. And this would give us only the *starting state* of achromatic vision, the receptor input, prior to retinal or cortical luminance processing.

That said, my own best guess is that the achromat will not lose whole snakes and walrus, or even medium-sized objects like bowls and buttons. Instead, the achromat will find it difficult to see *surface detail*. Again, examine the Gegenfurtner (2003) images particularly the S cone image (Fig. 4e). Of course, we don't expect oranges and apples to look *black* and this makes the fruit look rather peculiar in the S image. Strangeness aside, the fruit also appear somewhat plastic and featureless. Looking at the apple, its 'plasticity' results from the darkening of the apple image (by removing the 'red' light for the light source), an event that serves to highlight, by contrast, each white 'glint' of light from apple's surface. The apple appears 'plastic' because its surface looks too shiny to be real. However, the S-filter also makes it impossible to see the kinds of shading and shadowing that are so useful in normal trichromatic perception. Shading, which is the self-shadowing of an object given directional lighting, is a central cue for shape perception. For example, a round object, lit from the right, has both a circular boundary around its periphery plus a curved pattern of shading, from light to dark, on the left side of the image—two cues for a single property, roundness. We see surface texture by means of patterns of small shadows on the surface of an object: A dimpled orange peel produces a regular array of dimple-shaped shadows against the background of a bright orange surface. But if the image of orange itself is very dark, as it would be when filtered by an S filter, it would be very difficult to distinguish the dark dimples against their black background. (That is why a tight dress looks best in black at least for women with any surface 'texture'.) It is exactly this kind of contrast—patterns of low luminance contrast on reddish or blue-ish objects—that will be absent in rod achromatic vision. Hence surface texture and detail will be invisible on all but blue objects for an S-cone achromat.

In fact, this is why the designers at Carnovsky always wear black dresses. Well, not

really. Rather this is why the designers chose line drawings over photographs for the RGB wallpaper. In a line drawing, all of the surface features of an object, its texture and the shading from which we determine the shape of objects, is rendered with just line and empty space. By choosing three colours of ink with the same intensity, the surface features of the animals are rendered at a single level of intensity contrast, the difference between ink and paper. The designers could then heighten or ‘disappear’ this constant contrast through the judicious choice of light filters. But importantly, if a creature is visible in the RGB exhibit, so too are its features and surface texture—eyes, fur, feathers or scales. The careful choice of pigment and lights, plus the format of line drawings, accounts for the dramatic effects of the exhibit, the appearance and disappearance of whole creatures under different illuminants. But if natural images had been used, with a range of intensity contrasts typical of a high-resolution greyscale image, it is the surface detail that would have gone missing. A trichomat who views a set of natural images filtered by the rod sensitivity function might not find the loss very interesting. No creatures would suddenly vanish. But the cumulative loss of contrast information would amount to a (statistically) dramatic loss of information. *The rod achromat’s losses are no less real—and in all likelihood no less substantial—than the information loss that a trichomat suffers in the Carnovsky world.*

### **1.3 A Scenic Detour: Chromatic Processing**

In the last section, the central lesson was that the rod achromat does not see ‘in black and white’ if by that one means that the achromat has access to the intensity information represented by a black and white photograph. Of course, the achromat does not have access to wavelength/colour information but neither does the achromat have access the other dimension of a visual image, light intensity. Instead, an achromat has a very specific form of luminance information—image intensity filtered as function of wavelength by the rod photopigment, subsequently encoded as differences in photon catch, i.e. luminance contrast. And *that*, as should now be clear, is a different kettle of fish, the explicit consequences of the ‘empty’ definition of ‘luminance’ with which section 2.0 began.

In this section, I want to explain how chromatic systems arise and why, once in existence,

they turn out to be highly effective partners for luminance systems. On the one hand, the chromatic cells, which arise post-receptor in the retina, do not encode colour *per se*, neither the colour of objects and various media in the distal world nor what is called ‘image colour’, the wavelength composition of the retinal image and its spatial arrangement. Rather chromatic cells arise by chance, the inevitable outcome of the genetic variation and the wiring of existent luminance cells<sup>3</sup>. It is this chance composition that endows chromatic cells with highly complex informational properties, as opposed to the properties that a wavelength ‘detector’ would need to have. On the other hand, once chromatic cells are added to a luminance system, the informational ‘reach’ of the combined system greatly extends the informational reach of either component. Computational tasks that were beyond the capacity of either luminance cells *or* chromatic cells alone are made possible by the partnership of these two complementary systems. As the reader will have guessed, this partnership makes possible the perception of colour in the distal world, the colour of opaque surfaces, transparent media such as water, and of light itself. But as we will see in the next section, chromatic and luminance processing is also needed for the perception of lightness and darkness, to see coal as black and snow as white. Without chromatic encoding, the rod achromat lacks the veridical perception of surface lightness or darkness. In this second sense, in terms of darkness perception, the rod achromat does not see ‘in black in white’ either.

To see why chromatic and luminance systems make such good partners, back up a few steps. A central problem with a ‘pure’ luminance system is that it is rather primitive *qua* a source of information about the distal world. As light travels through the atmosphere, its interactions with bulk matter will affect, in specific and law-like ways, both its wavelength and intensity. For example, the intensity of light diminishes as it travels further and further from its source or as it makes contact with bulk matter; transparent substances act as wavelength filters thus changing the wavelength composition of the emerging light; opaque objects cast shadows, resulting in areas with diminished light intensity (shadows) but of the same wavelength composition. And so on through the laws of optics applied to a natural environment. Thus, each

---

<sup>3</sup> Note that receptors are neither luminance cell nor chromatic cells. Rather, the signals from various receptors are used as the inputs to luminance and chromatic cells, a distinction that will be explained more fully in a few pages.

dimension of the light stimulus, if encoded separately, would act *as an independent source of information about distal bulk matter*. By definition, a luminance system conflates these two dimensions of light. So a luminance system cannot make use of the full informational resources of the light *qua* multi-dimensional stimulus except serendipitously. So a luminance system alone is at a functional disadvantage relative to one that registers wavelength and intensity individually.

However handy it would be to have an independent encoding of wavelength and intensity, evolution does not strive *towards* anything. Physiological variations appear, they work or they don't. Once genetically entrenched, these chance variations tend to stay the course unless their retention results in positive harms. Now in our own case, for the evolution of our three cone photopigments, we can chart their paths in the lineage of Old World primates and the mammals that preceded them (Jacobs and Rowe 2004; Jacobs 2008; Jacobs 2009). For our purposes, the general story of how photoreceptors evolve (opposed to the specific evolutionary history of the three cones of Old World Primates) is the relevant one. Through genetic drift, random links in the amino acid chains of existing photopigments are altered. When a change occurs at a key location, a receptor with a new spectral sensitivity arises. In effect, existing photopigments mutate into new ones through chance substitutions. In each case, the result is a new spectral filter. Very occasionally, on the order of .01 duplications per gene per million years, an additional photopigment gene is created. And *that* photopigment will itself become subject to mutation and drift as time goes on. Over time, then, the existing photoreceptors of each species change their sensitivities, and on rare occasions, a species may gain an entirely new class of photoreceptor. As a direct result of the genetics of photoreception, each species 'auditions' a series of new photoreceptors, each with an individual sensitivity to light.

Above, in the section on luminance processing, I said that one of the central principles of vision is visual contrast: In order to see anything at all, the system must encode a difference in some property of light or another between two spatially adjacent areas of visual space. In the vertebrate retina, this requirement finds its expression in the centre-surround cells of the retina, the LGN and primary visual cortex. As its name suggests, a center-surround cell compares the total photon catch between two regions of visual space, between a central circular region of

visual space and the circular area immediately surrounding it (a sort of donut-shaped configuration). It is here, with the formation of centre-surround cells, that chromatic and luminance cells first emerge.

In a retina with exactly one kind of photoreceptor, a centre-surround cell will have—*can* have—only one configuration: it compares the absolute photon catch of the *centre* region of visual space with the total photon catch of the *surrounding* area. This is an *achromatic or luminance* cell by definition because it signals, for two distinct regions of space, the difference in photon catch (i.e. luminance contrast) *for one receptor/filter type*. One way to think of this arrangement—what a luminance cell does best—is that a luminance cell responds most strongly to any intensity changes that occur against a uniformly coloured background. In the terrestrial world of mammals, every visual scene contains numerous instances of this arrangement. It occurs whenever a shadow falls upon a uniformly coloured surface or whenever directional lighting produces shading. The spatial arrangement of a centre-surround cell insures that the background colour, whether it is green, blue or brown, makes no difference to the cell response. Because both the centre and the surround regions are ‘filtered’ by the same photopigment, both center and surround will react in the same way to the colour of background. The center and surround have the same ‘colour’ filters. So the colour of the stimulus is factored out and it is the intensity contrast that drives the cell response. A luminance cell is thus maximally sensitive to intensity contrast assuming a uniform background colour (and a colour within the spectral range of the luminance cell).

The addition of a new kind of photoreceptor to a retina—an event that is destined to happen regularly—makes for some interesting variations on this theme. The first option, the ‘if it worked once, why not try it again?’ option, is the creation of a new and distinct luminance system. These new centre-surround cells would signal the difference in photon catch, for the center and surround regions, by this new photoreceptor. This ‘same old’ option is actually more interesting than it first seems. If you look at the difference between the luminance images in Figures 4(c) and 4(e) for the S and L cones, you can see why this new luminance cell constitutes a different ‘take’ on the world. Because the S and L cones are different spectral filters, they will almost always differ in their total photon catch—and so too will their deliverances about

luminance contrast. Looking again at Figure 4(c) and 4(e), it is clear that the two types of luminance cells, centre-surround cells driven by either S or L receptors (but not both), would yield clearly different measures of luminance. A new photopigment equals a new luminance measure.

Note that the question “But which measure of luminance is *correct*?” is not a coherent question. Indeed, one might say that it misses the whole point of multiple luminance systems. As we know from the Carnovsky exhibit, each type of luminance cell is maximally sensitive only within a certain narrow range of wavelengths. So a new luminance system can extend the range of luminance contrast processing out beyond the spectral boundaries of a single photopigment already in place. Think of it this way. Fine-grained luminance processing is possible only within the small window of response for each photoreceptor. This is why a luminance system, with one receptor, results in such drastic information loss, relative to the intensity information of the image. As each new kind of luminance cell is added, the contrast range of the system as a whole is extended. At the limit, the system as a whole approaches a detector for *intensity contrast*. This is why multiple luminance systems are found in all diurnal mammals, to extend the range of fine-grained contrast encoding.

The addition of a second receptor also makes possible a quite different and equally interesting option. A center-surround cell could compare the photon catch of two *different filters types*, one for the centre region and one for the surround. This is a *chromatic* cell by definition a cell that compares the total photon catch of two different spectral filters or receptor populations. This new arrangement yields a cell with surprising properties.

Once again, look at Gegenfurtner’s S, M and L luminance images of common fruits and vegetables, here at the banana and the grapefruit (just under the banana and abutting it). Now imagine how the visual brain might go about the basic task of scene segmentation, of identifying the boundaries between objects, here between the grapefruit and the banana. In the L and M luminance images, the two fruits are a light grey; in the S image, the grapefruit is black and the banana is dark grey. To distinguish the banana from the grapefruit, any of the three images *could* be used: centre-surround luminance cells, fed by just one type of receptor for both the



centre and surround would respond to the banana/grapefruit boundary. But if you could compare the banana *in the L image* with the grapefruit *in the S image*, the difference in their luminance values would be striking. Instead of comparing two shades of grey (as in the L and M images), the banana (L image) and grapefruit (S image) have a boundary defined by black on one side, and white on the other, a high contrast boundary. This is why an S—L chromatic cell would be so effective (if it existed—this is a fictive example at least for human vision)(Fig. 8). Passing the cell over the border, it would compare the total photon catch of the S cones with the total photon catch of the L cones at the border between the two fruits. And *that* comparison would yield a very strong, highly reliable contrast signal—a *chromatic* contrast signal yet one demonstrated for us, quite admirably, with black and white photographs.

Importantly, the encoding of chromatic contrast does not depend upon any prior categorization of the distal surfaces or of the image areas into *colours*. A chromatic cell merely compares two different *luminance* measures. Nor does a chromatic cell *detect* wavelength contrast per se. For one, there are other stimuli apart from colour contrast to which chromatic cells respond (see below). But even so, a chromatic cell does not provide an objective measure of wavelength contrast. A yellow banana and a pinkish-yellow grapefruit are very similar vis-à-vis their surface spectral reflectance, the percentage of each wavelength of light that the two fruits reflect. So the colour (or more neutrally, ‘wavelength’) differences between a banana and a pink grapefruit are not very large. The response of a (fictive) S-L cell to their colour differences in the image, however, signals *high* chromatic contrast. An M-L cell, an equally fine example of a chromatic cell, would also respond to this boundary. But an M-L cell would yield a much lower chromatic signal. So a better way of thinking of a chromatic cell is as one that takes advantage of a fortuitous state of affairs, the difference in photon catch between two distinct spectral filters, one of which prefers short wavelength light, the other of which prefers long. At its best, a chromatic cell *highlights* a genuine wavelength contrast that exists in the world.

With the designation of ‘chromatic’ and ‘luminance’ cells in hand, however, it is easy to overlook the fact (given the human penchant for orderliness and simplicity) that both chromatic and luminance cells have complex informational properties. A chromatic cell is often called a ‘colour’ cell because it meets the formal neurophysiological definition of a ‘colour’ system, the

capacity to respond to wavelength independently of intensity information. Given two contiguous coloured stripes, a red stripe and a green stripe of equal lightness (intensity), an L-M ganglion cell will react to their common border. Hence a chromatic cell fulfills the formal definition of colour vision. Shine a flashlight on only the centre region of an L-M cell, however, and the L-M cell will produce a sustained response to this difference in light intensity. With a broad-band light stimulus (the flashlight) targeted on the center region, the center L receptors receive more ‘red’ light than the surround M receptors receive ‘green’ light (because the surround region is in the dark!). In other words, a chromatic cell does not *detect* wavelength differences but it will respond, vigorously, to certain wavelength differences. The same holds of luminance cells, *mutatis mutandis*. Within the range of greatest wavelength sensitivity for a luminance cell, it will respond vigorously to a difference in light intensity alone. Yet if one looks at the response curve for any photoreceptor, it is clear that particular, large differences in wavelength alone between two stimuli would make a difference in photon catch for receptors of the same type. So a difference in photon catch between a centre and surround, as registered by a luminance cell, *could* be the result of a wavelength difference alone. In sum, both chromatic and luminance cells respond to wavelength differences *of a certain kind* and intensity differences *of a certain kind*. Given their physiology, chromatic cells are more sensitive to wavelength contrast (along a particular colour axis) while luminance cells respond more actively to differences in pure intensity (within a certain spectral range). But neither a chromatic cell nor a luminance cell is specialized for the “if and only if” task of wavelength or intensity detection.

In sum, the essential difference between luminance and chromatic cells comes down to the filters involved in contrast processing — whether the comparison involves one filter or two. This way of describing the distinction represents it as primarily a distinction in ‘wiring’ or anatomy. In effect, I asked this question: If contrast processing is a hard constraint on the evolution of vertebrate vision, how many ways are there to wire-up a retina with more than one photoreceptor? And no matter how many more receptors are added, whether a species has two photoreceptors or ten, the answer is this: *Just two*. There are only two ways to make a luminance comparison: between filters of the *same* kind and filters of *different* kinds, luminance and chromatic cells respectively. Thus the distinction is one of ‘chance and wiring’, a division that occurs inevitably given the genetics of photoreception and rules of combinatorics. It is no

mystery, then, that the world contains both chromatic and luminance cells, that both forms of cells exist. Rather, the interesting questions concern the widespread integration of chromatic cells into the eyes all diurnal creatures and why the anatomy of biological vision respects this distinction throughout the visual system. If the species is diurnal and the environment contains light across a broad range of stimuli, it is overwhelmingly likely that the species will have both chromatic and luminance visual cells which comprise anatomically separate but often physiologically interactive systems. This universal phenomenon suggests that the luminance/chromatic divide has general informational consequences. But what would those be? What is it about chromatic and luminance systems that make them such good partners?

To answer this question, let us look at edge processing. Take a standard colour photograph. We can represent the separate contributions of wavelength and intensity to a colour photograph (or retinal image) with two separate illustrations, an isochromatic image and an isoluminant image (meaning, literally “all the same colour” and “all the same intensity”). Because light is a transverse wave and, by definition, every wave has amplitude and wavelength, every retinal image can be divided into these two components. So there are potentially two sources of information in a retinal image and at least on visual inspection they appear quite distinct (Figure 9). For example, in the isochromatic (or monochrome) image (Fig. 9b), object edges are demarcated by a difference in contrast with their backgrounds. As nature would have it, objects almost always differ from their backgrounds in intensity. In addition, the isochromatic photo also shows both shadows and object shading. In comparison, in the isoluminant image one sees only ‘colour without intensity’. Now although we often think of shadows as ‘coloured’ in the natural world, most shadows create only a difference in light intensity. Shadows on a green lawn create merely darker areas of (that same green) lawn. (Blue shadows on snow are the exception that proves the rule.) In the isoluminant image, the shadows and object shadings are no longer present (Fig.9c). What we see, in the isoluminant image, are object boundaries and expanses of surface colouring. When one looks at these two types of images, side-by-side, you can see that the object boundaries (and surface colour boundaries) are visible in both illustrations but the shadows are visible in only the monochrome image. Thus in a full colour image (Fig. 9a)—in the retinal image—objects (and surface markings) are demarcated by a combined edge of intensity and wavelength contrast, while shadows and shading are demarcated by intensity

contrast alone. *One way to differentiate object boundaries from shadow edges, then, would be to distinguish between the intensity and wavelength dimensions of the stimulus in the visual image.* Or at least, that would be a good way in principle. In practice, the tools at hand are chromatic and luminance cells not wavelength and intensity detectors. So the question is concerns which features of the world luminance and chromatic cells encode: what would a chromatic system paired with a luminance system make of the objective facts of object and shadow boundaries?

In an experiment by Hansen and Gegenfurtner (2009) they examined 700 images of the natural and artificial world to find out which edges, defined by either chromatic or luminance contrast, human vision can see. The image in Figure 10(a) depicts the original image or what they called the ‘Input Image’, a full colour image; the second image, 10(b), shows the ‘Luminance Image’ or the total photon catch by the system, calculated by adding together the absorption of the S, M and L cones; Figure 10(c) shows the M-L Image, namely the contrast between the responses of the L and M cones at each point in the image (using a continuum between red and grey to illustrate the relative responses). Finally 10(e) and 10(f) show the edges that can be computed from these two types of contrast data. Here you can see that the edges determined using luminance contrast were often distinct from—and different strengths than—the chromatic edges.

As the reader may already suspect, a test of a red-green (M-L) chromatic cell to determine what edges it can detect, is hardly likely to fail given an image of *red* fruit against *green* leaves. When the joint edge histograms for all 700 images were computed, however, the result was a general one. *Chromatic edges and luminance edges are statistically independent of one another in natural visual images.* So the addition of an M-L chromatic system to an M + L + S luminance system represents a huge leap in first-order information about edges: chromatic cells encode *the location and strength of edges in the image that luminance cells do not.* Equally importantly, a visual system with both a chromatic and luminance component gains *second-order information* about the *relative* locations and strengths of luminance and chromatic edges. With this comparison in hand, it is possible to differentiate objects from shadows using two simple rules: To find an object edge, look for a discontinuity that triggers both chromatic and luminance cells or for which the chromatic response predominates; to find a shadow, look for

contrasts to which *only* luminance cells (but not chromatic cells) respond. These are two rules that will work *most of the time* given the judicious selection, by evolution, of the different chromatic and luminance systems, and given a scene in line with the statistics of images of that species natural environment.

More generally, the addition of chromatic systems to luminance systems has proven useful because they are *complementary systems*. Chromatic systems respond vigorously to spectral discontinuities to which luminance systems have little response; luminance systems respond most vigorously to discontinuities in intensity in which chromatic cells are uninterested. In mammalian vision, this fact has been co-opted in the service of object vision. The anatomical segregation of the chromatic and luminance systems allows mammalian vision to utilize these two independent sources of contrast information *as independent*. Of course, scene segmentation through edge processing is one of the earliest and most essential capacities of any visual system. *But once luminance and chromatic edges are determined, they can be used, either together or independently, for virtually any visual task.*<sup>4</sup>

Exactly how luminance and chromatic systems work together (and apart) is something that recent experimentation is beginning to explore (Gegenfurtner and Kiper 1992; Cropper, Mullen et al. 1996; Baker, Boulton et al. 1998; Mullen, Beaudot et al. 2000; Mullen and Beaudot

---

<sup>4</sup> That said, some visual tasks *are* easier to solve using one source or the other—chromatic or luminance—information. For example, if the task is to track linear motion, a population of luminance cells (each cell with a ‘zippy’ onset and transient signal) will do this both more quickly and accurately than a population of chromatic cells (which have a sluggish onset and sustained response). But in the information game, *something*—some information—is almost always better than *nothing*. On a foggy, rainy day, you are outside your tent about to fry bacon in a pan (a typical day in Vancouver for any philosopher). A vaguely Grizzly-shaped object wanders into the middle distance. Under conditions of visual haze and scattering, luminance signals are often weak, noisy or often absent while chromatic signals are not as readily affected.. But it doesn’t matter that the more robust chromatic signal is a sub-optimal source of motion information. That the Grizzly bear has started to move (motion onset), that it is moving towards the viewer from a leftwards direction (direction of motion) and that it is rapidly increasing in velocity (2<sup>nd</sup> order velocity information) are bear-inhering-properties that it would be beneficial to perceive. The task is to use whatever information is ready-to-hand in the most optimal way, not to use the ideal source of information. This is where a chromatic system, paired with a luminance system, can provide clarifying information.

2002; Gegenfurtner 2003; Kingdom 2003; Kingdom 2005; Kingdom and Kasrai 2006; Kingdom, Wong et al. 2006; Gheorghiu and Kingdom 2007; Michna, Yoshizawa et al. 2007; Garcia-Suarez and Mullen 2010). Unfortunately, given the complex informational properties of luminance and chromatic cells/populations, predictions from first principles have limited utility without supporting models. Our best bet is to expect the unexpected. That is, we often assume, even when we know better, that the *point* of a chromatic or a luminance cell is to encode wavelength or intensity. We treat the informational complexity of chromatic and luminance cells as if this were a deficit or hindrance to be surmounted as soon as possible. (“Thank you God, for now the laws of optics are finally within reach!”) But the real story—the ‘there-is-no-simple-story’ story—is that the chromatic/luminance distinction was never ‘meant’ to track the dimensions of intensity and wavelength. It is a distinction of chance and wiring, as I said above, and one with which evolution has been grappling with since the advent of the first center-surround cell. We can be confident, I think, that primate vision stumbled upon any number of interesting ways to use the informational complexity of chromatic and luminance cells to visual advantage.

### **1.5. Albedo Perception: Perceiving Surfaces as Light or Dark.**

One of the cortical functions for which luminance information is almost certainly used is for seeing opaque object surfaces as light or dark. We see coal or briquettes as dark, copier paper and snow as light, and natural concrete as somewhere in between. As trichromats, of course, we also see these surfaces as having *colours*, what are known as the *achromatic colours*: briquettes are black, paper and snow are white, and untreated concrete is a medium grey. Given that one dimension of the trichromatic colours, both chromatic and achromatic, is darkness/lightness, there seems to be a relation between trichromatic *colour* perception and trichromatic *albedo* perception. Exactly what this relation might be, if any, is the subject of debate within the vision sciences. But the issue for the achromat, who does not see the colours, is much clearer. If the achromat sees surfaces as being light or dark, this must be the result of albedo perception proper.

The relevance of albedo perception to the current topic might well seem opaque to the reader: what does the perception of surface lightness by the achromat have to do with the

question of whether he or she sees the world ‘in black and white’? This very fact, that the relevance of albedo perception to an achromat’s visual phenomenology is not clear, is symptomatic of a deeper problem, a common misunderstanding about what is involved in seeing surfaces as light or dark. Let me take a moment, then, to discuss albedo perception before we wind our way back to the main point.

Jonathan Cohen, in a paper on colour properties (Cohen 2004), provides a good example of how we commonly think about albedo perception. In this paper Cohen’s main concern is to argue for a certain view about colour properties, the relationalist view. The topic of albedo perception arises only in passing in the context of his ‘Master Argument’ about colour constancy, how we decide the colours of objects in a complex scene containing directional illumination. Cohen’s example is a now famous photograph of a red coffee mug on a table, illuminated from the side by sunlight. Pointing first to the dark red handle of the mug (in shadow) and then to the bright cherry red body of the mug (lit by sunlight), Cohen asks how we choose the correct colour: Is the mug dark red or cherry red? Now, as Cohen notes, although the original photograph is in colour, readers with only a greyscale reproduction can ask a parallel question. If we examine the greyscale version, we will see a mug with a dark grey handle (in shadow) and a glossy light grey barrel (in sunlight), two distinct achromatic colours. Yet we still see the cup as having a uniform surface colour despite the variations of shadow and sunlight. We can thus ask a parallel question about *achromatic* colour: Which grey is the mug’s correct colour? Cohen poses the question of colour constancy, then, as a question about a shade of perceptual grey—almost black or light grey?

In computational vision, the problem of albedo is described somewhat differently. Below is a clear (and standard) definition (Anderson and Winawer 2005).

“The amount of light projected to the eyes (luminance) is determined by a number factors: the illumination that strikes visible surfaces, the proportion of light reflected from the surface and the amount of light absorbed, reflected and deflected by the prevailing atmospheric conditions (such as haze or other partially transparent media). Only one of these factors, the proportion of light reflected (lightness) is associated with an intrinsic property of surfaces and hence is of special interest to the visual system. To accurately recover lightness, the visual system must somehow disentangle the contributions of

surface reflectance from the illumination and atmospheric conditions in which it is embedded.” (pp. 79-80)

An illuminant (light) shines upon a three dimensional object. The intensity of light that falls upon each surface—the object’s *illuminance*—is a function of the brightness of the light source and the particular shape of the object. Each object, in virtue of its surface qualities, absorbs and reflects a certain percentage of the light cast by the light source, a property known as *surface reflectance*. Thus, the total light reflected from the object, its *radiance* is a function of both the intensity of light that shines on each point of the object’s surface and the percentage of that light which is absorbed. This light then travels to the retina and en route meets with certain media. Perhaps it is dispersed by smoke or haze in the air, transmitted through the coloured sunglasses or through the ordinary transparent lenses of correctional glasses. Even in the “normal” case, however, the light must travel through the atmosphere between the eye and object, then through the cornea and the lens plus the aqueous and vitreous humours of the eye. All of these media are filters, collectively known as *atmosphere*, that absorb, reflect and refract the light of the luminance image before it reaches the retina. Thus, conceptually, the problem of lightness perception for human vision concerns the disambiguation of the contributions of albedo from those other physical factors (illuminance and atmosphere) that result in the proximal stimulus, the retinal image (or luminance image in the parlance of computer science). In albedo research, however, researchers often choose to ignore atmosphere and concentrate on how to disambiguate the contribution of surface darkness from those of the light source. The diagrams in Figure 11 show this simplified problem in schematic form (Adelson 2000). Here, the three components of the computational problem are referred to as ‘images’ or layers, and the task is to disambiguate the three layers. In order to avoid terminological confusion, let us call what the schematic refers to as the ‘luminance image’ the ‘*retinal image*’.

Unfortunately, there is no easy solution to the problem of albedo perception even if one sets aside the vexed effects of atmosphere. Let  $L(x,y)$  be the luminance (retinal intensity) image,  $R(x,y)$  be the reflectance image, and  $E(x,y)$  be the illuminance image. Then:

$$L(x,y) = R(x,y) E(x,y)$$



Given only the luminance image alone, the problem is ill-posed or undecidable in the computational sense. For any value of  $L(x, y)$ , no unique value of  $R(x, y)$  can be computed without knowledge of the value of  $E(x, y)$ . Without  $E(x, y)$ , an indefinite number of images are possible. Exactly how the human visual system overcomes this problem is the subject of much debate. But for our purposes, the facts are not as important as the principles involved. What matters here is the general form of the various methodologies that have been proposed. Here, then, are two examples, followed by a summary of their shared assumptions

One prominent solution to the problem of albedo, proposed by Adelson (Adelson 1993) holds that the visual system uses various “tricks” and short-cuts based upon the properties of the retinal image. Adelson’s computations for lightness make use of the geometry of the image, how the light and dark areas of the image meet—what he calls “junctions”. Both the geometric form of the junction plus the relative luminance of each bounded area at a junction provides important clues about the causes of those edges. For example, Figure 12 (a) demonstrates how two identical edges could have two distinct causes in the world: the upper edge is caused by illumination differences while the lower edge is the result of differences in surface darkness. Figure 12(b) illustrates the various types of junctions used as clues in Adelson’s model— X, Y, L, T and junctions—and Figure 13 shows us how the junctions we can affect our interpretation of the scene. In Figure 13, the dotted rectangle, viewed alone, appears to be composed of stripes. But when the rectangle is viewed with one end or the other obscured (i.e. when we see different the contextual cues) the stripes within the rectangle are interpreted in two different ways. If you cover the right side of the illustration, the stripes will appear to be painted or seen the result of surface reflectance differences; if you cover the left side of the illustration, the dark stripes are appear to be shadows on the risers of steps. Adelson posits that we see the two sides as different because to the left of the rectangle, the junctions are arranged vertically, with their spines connected while to the right, the dark stripes form horizontal junctions. This arrangement of junctions combined with the arrangement of light and dark areas defined by their edges, determines our interpretations of the stripes. Adelson’s theory is considerably more complicated than this, but the above gives the reader the basic flavour of his view.

A second popular theory of albedo perception focuses upon properties of the scene that

are more directly tied to intensity contrast in the retinal image. For example, Knill and Kersten (1991) have demonstrated that shape information is critical to distinguishing whether an image area with a pattern of diminishing intensity is the result of illumination (on a curved surface) or surface reflectance (of a flat surface) (Fig. 14). More recently, Anderson (Anderson and Winawer 2005; Anderson and Winawer 2008), a leading proponent of this second view, demonstrated that the visual system appears to use both local cues (intensity contrast reversal across borders) and more global cues (occlusion and shape information) to separate a retinal image into “layers”, what are essentially depth planes corresponding to the plane of objects’ surfaces, the background and the foreground. In Figure 15, the disks on the left and right are physically identical. However we interpret the two figures very differently: The disks on left appear to be dark discs partly covered by a transparent white haze or fog; on the right, the discs look like white moons, obscured by dark clouds. When the disks are turned the 90 degrees, however, much of the illusion is lost. Anderson argues that the rotation destroys the pattern of contrast reversal at the edges of the disks that is necessary for dividing the image into depth planes. Without depth information, there is no clear disambiguation of surface reflectance (light or dark surface) and into the atmospheric factors (fog or dark clouds). We also lose the strong illusion that the two sets of identical disks are entirely different. Insofar as the illusion remains, this is attributable to the overall illumination of the two backgrounds, one of which, as a whole, is statistically darker than the disks, the other of which is lighter. On Anderson’s view, then, intensity contrast across a figure-ground border is essential to our perception of lightness and darkness as are the global scene characteristics. The sort of local and global ‘clues’ that Anderson suggests are of a different kind than Adelson’s image junctions.

The reason for the above comparison was to give a sense of how the problem is conceived in computer science and what would constitute a solution relative to that conception. Clearly, there is significant disagreement among computer scientists about how our visual system deals with what is a theoretically intractable problem: each posits different sets of cues and, as a result, each theory would predict different occasional failures and illusions of albedo perception ((Adelson 2000; Corney and Lotto 2007; Poirier 2009; Anderson and Khang 2010; Spillmann, Hardy et al. 2010)). On the other hand, there is also significant agreement about the nature of the problem and hence about what a system for albedo system must do. All researchers

agree that:

1. Luminance, illuminance and surface reflectance are physical properties.
2. Surface reflectance (lightness) is a constant property of object surfaces and hence is useful for identifying objects, distinguishing one object from another, and tracking them.
3. In systems of natural vision, albedo perception is a computational process that uses intensity contrast data from the retinal image to determine surface reflectance values. (Intensity contrast could be used to see different kinds of junctions or it could be used to see more large-scale patterns in luminance contrast itself.) The result in the perception of an object surface *as* having a specific lightness or darkness, an intentional property of the object.
4. Albedo perception for surfaces within complex natural scenes depends upon numerous cues within the image. As the number of natural cues diminishes—that is, as the scene/luminance image becomes less complex—mistakes in lightness perception, *lightness illusions*, occur more frequently.
5. Given that there is no decidable procedure for solving the problem of albedo perception, human lightness perception is not 100% reliable. Yet through the use of multiple cues, human lightness perception approaches a reasonable standard of accuracy.

Turning back to central problem of this paper, the nature of achromatic experience, let me make explicit why this conception of the problem of albedo is at odds with what I'll call, for the lack of a better name, Cohen's Conception. Going back to Cohen's example of the coffee mug in shadow, observers agree that the coffee mug is a uniform achromatic colour despite the numerous greys caused by the directional illumination. Cohen then asks which shade of grey is the true (achromatic) colour, a question that seems to make sense. But if one goes back to Adelson's three figures that illustrate the problem, Cohen's question appears suspect. The first illustration, what we are calling the Retinal Image, is a representation of the intensity values of light at each point in the image reflected from the distal scene. In the second illustration, the Illuminance image represents the effects of a directional light source on a block. Finally, the

Reflectance Image represents the surface reflectance of that same block, the constant proportion of light reflected from the four cubes of the block, the result of its inherent surface properties. Thus in each representation of the block, the greyscale is used to 'stand in for' a different property: image intensity, illuminance of the object, and surface reflectance. Suppose then, that using these illustrations, we point to two different locations in the Retinal Image, say the endpoints of the arrows p and q and we ask 'which shade of grey, medium grey or dark grey, is the small cube *really*?' In the Retinal Image, those two different greys, indicated by the arrows, represent different values of *light intensity*; in the Reflectance Image, the single shade of grey of the uppermost right cube represents a dispositional property of the cube's surface, *surface reflectance*. Neither shade of grey 'is' the correct albedo of the small cube: the two greys do not represent albedo. The question is incoherent.

Cohen's original question about the correct shade of grey initially 'plays' much better because of the dual nature of photographs. The photograph of the cup is an object which itself has albedo, the many variegated greys that make up its surface; the photograph *qua* representation of a scene also portrays a coffee cup, a cup which has a certain surface lightness or darkness. When we are asked whether the cup is 'this grey' or 'that grey', we are looking at two areas of the photograph with different albedo, one near black and the other light grey. Thus it seems to make sense to ask *of the cup* which of the two greys 'matches' the cup's surface. But this is a false dichotomy. There is no reason why the surface reflectance of the *photograph*, at any point, must be the same as the surface reflectance of the *mug* it represents. If the coffee mug had been photographed entirely in dark shadow, its photographic image would have been a montage of very dark greys. Yet we would still *see* the cup as medium grey.

A few more general conclusions can now be drawn. On Cohen's view, the problem of albedo perception is primarily a question of phenomenology: Is the mug *this* grey or *that* grey? As such, it is problem in the first person: the observer must select the correct phenomenal grey. This question seems to make sense as we view the photograph before us with its stable properties of surface albedo. The problem of albedo, as construed by the vision sciences, is the computational question of how a visual system disambiguates two properties of the visual scene conflated within the retinal image, surface reflectance and the effects of the light source upon

that surface. Thus construed, the albedo problem is a question of sub-personal processing, of how to account for what we as observers see, namely surfaces as light or dark. It is the essential question of how we gain intentional perceptual representations of the distal world, of how we see the multiple properties of the world from the first person point of view. On this view, the question ‘how do we see the lump of coal as dark?’ is the same type of question as ‘how do we see the chair *as* being behind the table?’ (a question about depth perception) or ‘how do we recognize a certain grapheme *as* being an ‘F’? (a question about object recognition). For the computationalist, then, our very ability to ask the question posed by Cohen—‘which is the correct grey?’—is a question that can be asked only *after* all the hard work of viewing a natural scene is done. We gaze at the photograph of the mug on the kitchen table and interpret it as we would the Retinal Image of a natural scene: we see a mug with a uniform surface lightness with its handle in shadow, its body illuminated by bright directional light. We can then examine the photograph and determine *its* albedo qua surface of a paper photograph, there discounting any effects of the illuminant in the actual world, the one in which you, the observer, examines the photograph (does your shoulder cast a shadow over the image in front of you?). With these two processes of albedo perception behind us, we can then ask Cohen’s question about achromatic colour—‘which is the correct grey?’ Looking at the photograph before you, you can ask this first person question only after two problems of albedo perception—for both the mug and its medium of representation—have already been finessed.<sup>5</sup>

This way of thinking about albedo perception may seem a bit odd until one remembers that light itself is a perceived property of any distal scene. We see the dark shadows in the forest (this is what frightens us), white ‘glints’ of sunlight off water (why we reach for sunglasses), and

---

<sup>5</sup> I haven’t discussed how the brain figures out a sensitive measure of absolute surface reflectance as opposed to an assessment of a relative lightness and darkness within a given scene—e.g. the coal is much darker than the table on which it sits. I am myself somewhat skeptical about the utility of such a representational process, whether the brain bothers to assess lightness or darkness in the relevantly fine-grained way that would be necessary to assign each surface a place on an absolute scale from dark to light. Nonetheless I believe that we think of lightness and darkness as an objective property, a continuum along which each surface has an objective place, whether or not our on-line assessment of albedo is, in practice, merely ‘good enough’ for whatever task is at hand. Whatever the answer to this question, however, albedo perception does not come down to choosing the right grey ‘chip’. It’s not a question of phenomenology.

the fiery sunset reflected in a wall of windows (why we reach for our cameras). Given a retinal image, we could not see the distal scene unless, sub-personally, the visual system was able to distinguish properties of illumination from properties of objects (Arend and Spehar 1993; Kingdom 2008; Kingdom 2011). This is why the problem of albedo could not be a question of phenomenology *per se*, of comparing, from the first person, the different grays in an image. One of the first and most important tasks of mammalian vision is the division of the scene into illumination and object properties, a task that must occur *prior to* our intentional perceptions of objects and illuminant properties. For example, to see a cube as a cube in a scene with directional lighting, there must be a way to reconcile the proposed object shape, being cube-shaped, and the pattern of illumination. Such a reconciliation would not involve a *precise calculation* of a measure of *absolute albedo* for each visible surface. But it would involve *rough and ready* assumptions about *relative albedo*, e.g. that all of the visible sides have roughly the same surface reflectance, or that the front-facing facet must be lighter than the top surface in order for the object to be cube-shaped. In other words, when, from the first person, we peer at the scene (or image) before us, it is part of the content of our perception that objects have a certain shape consistent with the conditions of directional lighting, perceptions that presuppose at least some assumptions about relative surface lightness (e.g. that this surface is darker or lighter than that one).

Finally, we can return to the achromat. If we think of the problem of albedo perception following Cohen, as a choice between two phenomenal greys, then an achromat, who sees coal as dark, also chooses among some set of phenomenal greys. Thus there is no real question about an achromat's visual experience as a result of albedo perception: it *is* like looking at a monochrome photograph. *It is the same, by definition.* But if we treat the problem of albedo perception as a computational conundrum, then there is a substantive question to be asked about the achromat and his or her experience: Is the achromat capable of seeing surface reflectance given the luminance information of the rod achromatic vision? Note that the computer scientist starts with a bird in hand, an intensity image of the distal scene. Any biological system for object vision starts with a luminance image that, as we have seen, can be of many types. It is this fact that puts the achromat at a true disadvantage for albedo processing—a rod-only luminance system. Without wavelength information, the achromat cannot compute the intensity values of

the retinal image and with only one photoreceptor the achromat cannot even approximate image intensity (by combining multiple receptor outputs). So the achromat has no chance to regain this intensity information at any stage of vision.

It is worth pausing to consider the extent to which this puts the achromat at a disadvantage for lightness perception. If one were to give the achromat seven or eight paint chips of different colours side by side, each of which was equally bright, the achromat could not see that this was so. Each would appear to have a different lightness value. Presented with a series of coloured paint chips of varying brightness, an achromat could not order them from light to dark. With only a single receptor, the rods, with a spectral sensitivity centered in ‘green’ range of light, the reflections of red and blue objects will be very dark, the reflections of green objects, very light. An achromat might be told that his favorite tie is composed of bright magenta and turquoise stripes but the achromat will see neither the pink nor the blue stripes as bright. His distinctly cheerful tie will appear as somber attire to the achromat, a tie with an overall pattern of low contrast dark stripes—a good tie to wear to, say, a job interview or, better, a funeral.

So the rod achromat suffers a large deficit in lightness perception. However bad at the task the rod achromat may be, he or she nonetheless perceives the surfaces of opaque objects as having a *constant* property. Suppose you were asked to view the world using a virtual reality mask. The camera, from which your mask receives images, uses a narrow-band green filter and it transmits this information in the form of greyscale images. The camera is pointed at a real table covered in blocks of many colours and your task to arrange them in various ways. With no knowledge of the actual colours of the objects, you would see each cube as being light or dark. If told to build a tower, you would be able to stack and move the cubes using this constant surface property even though, were you to think of the blocks as light or dark, you would be almost certainly wrong in your perceptual judgments. This is the position of the achromat. He sees surfaces *as* having a constant property, ‘pseudo albedo’, which is distinguished from the effects of any particular illuminant on the scene; he knows that this property is not what others would call ‘lightness’ or ‘darkness’, and; he knows that his judgments of surface lightness will usually be wrong. There is a very robust sense, then, in which the achromat’s perception of

albedo is relevant to achromatic phenomenology and in which it widens the gap between the achromat's and the trichromat's visual phenomenology.

## 5.0. Conclusions

This paper began with the claim that the common division between black and white and colour leads us astray when we think about the nature of human visual phenomenology and the neural processes that support it. The running example has been that of the achromat whom Nordby claims sees in black and white. My reconstruction of that argument was based on the commonalities between three types of experience: the prototypical case of 'seeing in black and white' in which the trichromat experiences looking at a black and white photograph; the trichromat's visual experience at night when only the rod luminance system is active, and finally; the visual experience of the rod achromat who has only a system for night vision, the rods and the magnocellular system. By transitivity, if the trichromat sees in black and white, so too does the rod achromat. I expect that most readers will now see some of the flaws in this argument but let me walk through the three different cases in order to make explicit their differences and commonalities. The question at issue is whether there is a sufficient overlap, in physiology, to justify the conclusion that the achromat sees 'in black and white'.

Let's start first with the trichromat who views a black and white photograph. When a trichromat looks at, say, Arthur Sasse's famous portrait of Einstein (with protruded tongue), she looks at a physical object that reflects light continuously across the spectrum of visible light. This object also meets the following condition: *either* it reflects each wavelength of light equally, relative to a set intensity of light, *or* the photograph reflects light such that, for any given point in the image, the photon catch of the three cones is the same. This strange disjunction exists because a light wave has both wavelength and amplitude. So an image that conveys intensity information necessarily reflects light of some wavelength or other. The convention of black and white photography attempts to render the image such that the light it reflects has no discernible predominant wavelength—discernible by the trichromat. A black and white illustration can be printed with black pigment/ink or with the three standard colours for printing, cyan, magenta and



yellow. If the three inks are each combined in the correct intensities, the three colour ‘filters’, the S, M and L cones, absorb photons in equal measure at each point in the photograph.

The portrait of Einstein, viewed by a trichromat, produces a reaction in all three cones and thus produces signals in visual channels of all types, luminance and chromatic. In the luminance channels, each luminance contrast cell will respond to *only* intensity differences. There are no wavelength differences in the photograph, and thus any differences in photon catch between the center and surround area of a luminance cell indicates intensity differences. Now, in the trichromat, there will be many different luminance channels, each of which uses a distinct spectral filter. In the normal case, when the trichromat views the world, each type of luminance cell would produce a different measure of luminance contrast (for luminance is relative to a filter type). Here, though, light reflected from Einstein’s photograph (as opposed to Einstein) produces the same absolute photon catch in each spectral filter. So all luminance contrast cells will produce the same signal, an *intensity* signal, relative to any point in the visual image. That is the luminance side of the equation. In the two chromatic channels of the trichromat, the chromatic contrast cells, with input from two different spectral filters, will *not* ‘highlight’ wavelength contrast: There is *no* wavelength contrast in the portrait. In effect, a black and white photograph does not silence chromatic processing. Rather it neuters it. The signals of chromatic cells, of whatever type, also indicate stimulus intensity because each type of cone responds with the same photon catch *by definition* (of a black and white photograph). The upshot is that all contrast cells, luminance and chromatic alike, encode intensity contrast across a spatial border—an unsurprising result given that this is the function of modern black and white photographs. It is the recognition of this fact, at some higher level of visual processing, which *may* give rise to the perception of the photograph *as* black and white.

When the trichromat sees the photograph of Einstein *as* a black and white photograph, this would normally involve two different types of intentional perception, of the photograph *as* having only *achromatic colours* (the colours from black to white) and *as* having areas that are *light and dark*. If you go to the paint store and choose ‘Ripe Aubergine’ as the new colour for your dining room, you are well aware that the colour is a dark one. That is what will make it a cozier and more intimate space. But you do not think that ‘Ripe Aubergine’ is an *achromatic*

colour, a shade of grey. At least in some circumstances, the perception of surface chromatic *colour* and the perception of *albedo* come apart and so too must the processes that produce these perceptions.

Suppose then the trichromat sees the portrait of Einstein as having light and dark areas. Even though a black and white photograph produces, in the retina of the trichromat, a unique ‘signature’ of retinal encodings (i.e. the ‘neutering’ of the chromatic systems, etc.), albedo perception nonetheless requires higher-level processing. (This is why I said, in the concluding sentence of the paragraph before the last, that such retinal signals *may* result in albedo perception.) When we look at the portrait of Einstein, the chromatic and luminance responses of the trichromatic retina indicate that there is neither a predominant wavelength nor any wavelength contrast in the retinal image. But albedo perception is the perception of the surface lightness of distal objects, not of the retinal image itself. This is the difficult part to take on board: Albedo perception, even for a black and white photograph, requires a complex computational process. The very same considerations that apply to seeing a 3D block figure, like the one in Fig. 11, apply to Einstein’s portrait. For example, if the photograph faces towards the light source, it will be brightly illuminated; if the light source is directly behind it, the portrait will be darkened. In addition, nearby objects, such as your own body, can shadow the portrait in any number of ways. Any of these conditions will change the retinal image, reduce or increase its brightness overall, or have a selective effect on the intensity of its various parts. So the problem of albedo still stands: in computational parlance, the observer must separate the Illuminance Image from the Reflectance Image of Einstein’s portrait, given the retinal image. During this process of albedo perception, the ‘signature’ chromatic and luminance signals produced by the black and white portrait will be used to infer the complex pattern of light and dark areas that defines its surface. But the albedo perception is not a process of phenomenal reconstruction—a systematic mapping of the intensity values of the retinal image into a visual area of the phenomenal greys. For one, such a process would not yield a veridical perception of the albedo of the photograph. For another, the process of albedo perception is one of complex *inference* not one of *selection*, in which the observer, from the first person, selects the correct shades of grey from a reproduction, in phenomenal greys, of the retinal image.

In the usual case, the trichromat will also see the Einstein photograph *as* black and white, as containing only achromatic colours. Clearly, a discussion of the perception of colour, achromatic or otherwise, is beyond the scope of this paper. But I think that one point can be made here, given the discussion that has come before: The perception of achromatic colours is unlikely to rely upon entirely low-level visual process. This is why trichromatic observers are often *wrong* about black and white photographs, why we sometimes see a subtly tinted photograph as containing as black and white and why we sometimes see a black and white photograph as containing chromatic colours. For example, there is a contemporary Slovak photographer, Peter Župnik, who applies very faint pigment to black and white photographs, often to the lighting within the scene as opposed its objects. Watching a person examine a Župnik photograph is very interesting: the viewer leans into the photograph, scanning it repeatedly, attempting to discern whether the photograph is black and white or ‘coloured’, a surprisingly difficult task for most of Župnik’s photographs.<sup>6</sup>

The reverse effect—seeing a black and white photograph as coloured—can occur as well. In a now famous experiment, Gegenfurtner and colleagues (Olkkonen, Hansen et al. 2008) showed subjects a series of photographs, one after the other, each containing the same image of a banana. In the first image, the banana is coloured a saturated blue; each successive photograph involves a change in hue along the blue/yellow colour axis; in the last image, the banana is coloured a saturated yellow. A single frame, in the middle of the series, contains the null point, when the image is neither blue nor yellow but entirely grey. Subjects were told to press a button when the banana first appears yellow. Invariably, subjects choose the null point, the grey image, as the first yellow image. So our knowledge of and memory for surface colour can affect our current experience of a black and white image. Note that in both examples given, of the Župnik photographs that appear to be black and white (but are not) or of the banana that appears to be yellow (but is not), the problem is not one of *illusion*. An indefinite number of photographs could have been used in the Gegenfurtner experiments as long as they depicted prototypically coloured objects. Nor is the colouring of the Župnik photographs too subtle for trichromats to see. On the

---

<sup>6</sup> You can view Zupnik’s photographs at <http://www.zupnik.eu>. The photograph in my dining room is one of the Day’s Dreams series, ‘The Private Investigation’. In this photograph, the steam rising off the pig (yes, it is a pig) is tinted blue—the colour of night, of course.

contrary. The cognitive problem is that we both expect to see colour (given daily vision) and expect not to see colour (in a black and white photograph), and this makes it very difficult to discern *what* we are actually seeing. It would seem that whether we see an image as black and white or as coloured is also a matter of sophisticated perceptual *inference*, not a function of the brain's information that wavelength contrast is absent in primary visual cortex.

Finally, the normal trichromatic sees the *objects represented* in the black and white photograph, here Einstein and his tongue, as light or dark. This is albedo perception applied to the objects of representation. With intensity information in hand, adult trichromats are seasoned interpreters of visual images including the surface lightness of represented objects. We trichromats will not be fooled by an image of a white egg sitting in dappled shadow. We will not see a dark egg with darker leopard spots, but a uniformly white egg in dappled shadow. (Indeed, even if we viewed the same image under a piece of translucent red film, we would still see the egg as white and the dapples as a feature of illumination (as long as we could see the red film as a transparent colour filter.) On the other hand, adult trichromats do have trouble 'toggling' between the lightness/darkness of objects as depicted and of the surface properties of the representation itself. This fact is nicely illustrated by Adelson's Checker-shadow Illusion (Fig. 16). In the Checker-Shadow Illusion a green cylinder is depicted as sitting on a black and white checkerboard, seemingly lit by a light source on the right (the cylinder 'casts' a shadow over the checkerboard to its left). When we look at the illustration we see the checkerboard as one normally does, composed of a pattern of black and white squares. But if we are asked about the albedo of the illustration itself, that is more difficult. There two squares, A and B, that are represented as black and white which in fact have the same albedo in the illustration—the white square in shadow and the black square that is (represented as) fully illuminated. If you place a mask over the image that reveals only the squares in question while obscuring the rest of the scene, it is easy to see that these two areas, A and B, of the illustration have exactly the same surface lightness. We cannot see the two areas of the illustration as they are, as opposed to how they are represented as being: This information is lost to the first person point of view in the very process of the interpretation of the scene. It is not surprise, then, that when the scene is obscured we can have a veridical perception of the areas A and B qua surfaces of the illustration itself. Presumably, the infant must learn to see represented albedo while the adult trichromat

must try use visual aids *not* to see it, once the capacity for albedo is learned.

What happens to the trichromat when he or she views the world at night? Perhaps the most important difference between night vision and the perception of an achromatic photograph is that in night vision, the trichromat views *the natural world*, a world of coloured objects, not an artificial image designed to convey intensity information. We do not see the colours of objects in the dark, but whether the world is bathed in sunlight or moonlight, all surfaces reflect and absorb light as a function of wavelength and intensity. And of course, on the receiver end, rod luminance is still a function of the intensity and wavelength of the stimulus, with a spectral response very much like the Carnovksy Green Monochromat. So, the reflected light of red and blue objects will result in low photon catch; green objects will cause high absorption. In other words, colour matters even at night. At the level of the retina, the rod achromat and the dark-adapted trichromat have the same visual access to the world.

What makes trichromatic night vision interesting, and different from rod achromatic vision, is trichromatic visual memory, both personal and sub-personal, of object properties as seen in daylight, i.e. trichromatic visual knowledge. When the dusk falls on the trichromatic system, the photon absorption of the cones will gradually lessen until their catch is so low that it becomes impossible to distinguish signal from noise. Below a certain level of illumination, chromatic cells respond with idle or random chatter. (Remember that the rods will feed into the magnocellular luminance pathway, thus creating a luminance signal at night.) Tellingly, this is not the ‘state of the union’ of the trichromat who views a black and white photograph. In that case, the chromatic and luminance cells—all of them—detect *contrast*. A comparison of these different signals would yield the conclusion that the proximal stimulus, the retinal image, does not have (discernible) wavelength contrast or a predominant wavelength. It is this information that allows the trichromat to see the photograph as ‘black and white’ in either sense (as having only achromatic colours or containing a pattern of light and dark areas) in daylight. But at night, no such information is available. Rather, the trichromatic visual system registers that the robust luminance signals are accompanied by a marked *absence* of any signal from the chromatic pathways. With this information in hand, the trichromatic infant will learn *that the world is dark*, the most likely conclusion consistent with the signal pattern and with the overall reduction

in illumination. (It is also possible that the lights have gone out *and* that world's surfaces have simultaneously changed to achromatic colours but this is less likely, as an infant will learn.)

So how does the trichromat see surfaces given visual memory? The trichromat, like the rod achromat, cannot perceive albedo from the rod luminance signals alone. Like the rod achromat, the incoming signals provide the basis for seeing only the strange property of 'pseudo albedo', a measure of surface *luminance*, not surface *reflectance*. These perceptions of pseudo-albedo are inconsistent with the trichromat's memories of object surfaces under daylight. They are also inconsistent with her memories of familiar and prototypically coloured objects—e.g. the cherry blossoms as pale pink as opposed to a dark colour of some kind. Here we arrive at *terra incognita*: we do not know whether or how past experience of albedo and colour influences trichromatic *night* vision. Does the trichromat merely discount her daytime perceptions of albedo and colour when she views a familiar scene at night? Is the trichromat's belief that we cannot see colours at night extended to her perceptions of albedo? Or does daytime knowledge of albedo and colour taint ('tint'?) her nighttime perceptions? This is a question that would reward careful investigation. Certainly if you ask a class of students whether they find their cars by colour at night in a parking arcade illuminated by low energy sodium lights (which produce only 'orange' light), many students will claim that they do, that they can see colours under 'orange' light. It is also clear that our perceptions of the world—even colour perceptions—can be influenced by our memories. The Gegenfurtner experiments (Olkkonen, Hansen et al. 2008) with the blue/yellow images of the banana suggest that prototypical colours influence our perception of an achromatic photograph. Trichromats are also well aware of which properties of the world are stable and which are not. As the light grows dimmer at dusk, we become less and less able to see the colours of the world. But we do not see the colours themselves disappearing, chromatic colours turning into their achromatic cousins, much less previously light or dark objects assuming a new albedo as Nordby so picturesquely suggest ("certainly (all trichromats) must have witnessed the gradual disappearance of colours when darkness sets in"). But this is not the experience of the trichromat: we are aware that darkness hinders visual perception. Our knowledge of the visual world cuts both ways: An adult trichromat knows the colours of familiar and prototypical objects and that these relatively stable properties may be difficult to see at night. What a trichromat *does* see at night is an open question

Finally, we can return to the rod achromat. Much of this paper has been concerned with showing just how different the rod achromat's experience must be relative to that of the normal trichromat. For example, a trichromat has a broad range of visible light and hence a broad range of fine-grained luminance contrast information. A rod monochromat, with but one receptor, has a narrow spectral range of visible light and fine-grained luminance contrast is confined to that narrow window. The rod achromat sees the world 'through a glass darkly', with less light given only one receptor and with spectrally biased light in the bargain. For another, the trichromat has a multiplicity of chromatic and luminance channels. So the trichromat is sensitive to edges within natural images, the chromatic edges, to which the achromat has no access. Through the use of both chromatic and illuminance contrast, the trichromat can perceive both the achromatic and chromatic colours, and the inherent lightness of opaque surfaces. The trichromat can also use these types of information as independent sources of knowledge of edges in the world, a fact that opens the doors to numerous new processing strategies. The achromat, in contrast, cannot make an accurate judgment about surface lightness and cannot see any colours at all, chromatic or achromatic. Finally, the trichromat enjoys excellent spatial resolution in daylight conditions, and hence has the kind of high acuity information necessary for depth perception by stereopsis, the perception fine-grained surface patterns, for the control of fine motor skills, and so on. The achromat never, under any conditions, achieves comparable spatial information. It is clear that the achromat suffers a general impoverishment of information compared to the trichromat.

Now, *prima facie*, these differences may not seem relevant to the problem at hand. Nordby's argument assimilates achromatic visual experience with trichromatic *night* vision not with normal daytime vision in the trichromat. Surely, one thinks, Nordby is on safe ground here. But the background assumption that the trichromat and achromat have the 'same' luminance system, based on the anatomy of the rods and the magnocellular pathway, holds only if human visual development follows a fixed and unalterable path, if there is no plasticity in the development of vision based upon experience. And this, we know, is false. Visual development in mammals is thought to proceed via a developmental cascade of neural function. Incoming information in the first days or months after birth determines the nature of low-level visual mechanisms, the cells responses and their topographical organization in the LGN and V1. Each

subsequent step depends upon the system's current state (hence on past input) and upon current incoming signals. This developmental cascade—the sequential extension of visual function in response to input—continues until maturity. Thus both the type and timing of information determine a system's end state. It also follows that the trajectories of two systems, however similar they may be at the start, can diverge during the course of development as a result of asymmetric input.

In this case, the visual development of the human trichromat and achromat, we know that the visual input is very different from the first moments after birth onwards. Although the human M+L luminance system develops fastest, followed by the S-(M+L) system and then the M-L chromatic channel, by 5 months of age the trichromatic infant is approaching an adult-like sensitivity to contrast in all channels. So with the exception of the first few weeks post birth, a trichromatic infant has access to achromatic and chromatic signals of many kinds. We also know that the learning conditions for vision favor the trichromatic infant. Assuming that most visual learning occurs when the infant is awake, in the presence of an adult trichromat who prefers light to dark conditions, the visual input of trichromatic infants will occur under optimal visual conditions for the trichromat (i.e. bright light) while achromatic infants must learn under the worst possible visual conditions for rod vision (i.e. bright light).

Finally, one of the central arguments in this paper was that adding chromatic function extends the informational reach of any visual system. Chromatic cells highlight one range of contrasts within natural images, luminance cells capture a different one. The two types of cells, and their various sub-types, tap into independent sources of information about edges. It is this difference that makes chromatic and luminance systems such excellent partners. So the presence of multiple chromatic and luminance signals in the trichromat makes possible, both logically and empirically, *different strategies* for the representation of distal properties. This fact, that there are two independent sources of information physically shapes the trichromatic visual system to accommodate the requisite parallel and/or joint processing of these signals. In other words, the dual encodings both make possible different strategies *and* change the physiology of human vision as it develops. It is thus very unlikely that achromat and the trichromat will have magnocellular systems with the same functional capacities. Without chromatic input, the



achromatic system must either solve the visual problems differently, using only narrow band luminance input, or not all. Indeed, we can expect the achromat's magnocellular channel to be specialized for the processing of rod signals alone, with the all of the attendant differences in information of rod vision. Indeed it may be the case that while trichromat's have a visual huge advantage under daylight conditions, the very innovations for chromatic and luminance interaction, may turn out to be a hindrance to night vision. This is not unlikely, that two developmental cascades with distinct resources, may have specialized in divergent ways. In any event, it seems unlikely in retrospect that the trichromat and achromat 'share' a luminance system in any substantive sense.

We come then, back to the initial question of this paper: Is there any basis for saying that an achromat sees 'in black and white'. The reader will have noticed that I did not, at the outset, offer a definition of this central term. I chose to simply adopt the prototypical case of seeing in black and white and then looked for an overlap in physiology between the prototypical case and the rod achromatic vision that would explain why both observer's had a shared phenomenology. At this juncture, there would seem to be very few options for this common property. 'Seeing in black and white' cannot be matter of viewing a scene on the basis of luminance information alone. A trichromat uses chromatic information to see a black and white photograph as black and white. So the prototypical case does not meet this condition. Or perhaps the suggestion is that the luminance content alone explains their shared phenomenology. Here, though an argument would be needed to explain in what sense precisely an achromat and a trichromat have systems with a shared content. Recall that luminance is a measure of information that is always *system relative*. If you have a luminance system that depends upon an S cone and I have luminance system that depends upon an L cone, in what sense do our perceptions have 'the same' content? One cannot claim that both systems encode *intensity* information about the scene for they do not. Yet there has to be more to this claim of common content than the bare fact that both systems access the world via a single (but different) receptor suited for photon absorption under daylight conditions. What exactly is meant by the same content? Second, seeing in black and white cannot be a matter seeing a scene as having only achromatic colours, or seeing a scene as being devoid of wavelength differences or a predominant wavelength. All of these abilities require chromatic information that the achromat does not have. So again, seeing the achromatic

colours could not be requisite common property.

Finally, we come to albedo perception, a capacity that both types of observer *do* have in common at least if one includes the ‘pseudo-albedo’ perception of the achromat. Of all the visual capacities that an achromat and trichromat might share, this seems the most likely. But would a shared capacity for the perception of surface lightness account for the wholesale phenomenology of seeing in black and white? If you think of albedo perception à la Cohen, as choosing the appropriate shade of grey for each surface in a scene, this suggestion makes intuitive sense. Once each surface (or translucent body) is so coloured—or at least achromatically coloured—everything that can be seen now has some greyscale colour. The whole world is now coloured in black and white. And *that* would seem to be a good candidate for the phenomenology in question.

Albedo perception as described above, however, is not a matter of choosing the correct phenomenal grey and projecting it upon each surface/interior of each opaque/translucent object. Like shape perception, albedo perception is the perception of a property based upon a complex computational or other neural process. It begins with low-level luminance contrast information as its input and ends with a systematic representation of surface reflectance independently of atmosphere. It is the ability to see lightness, e.g. to see whether the egg is white independently of whether the egg is in dappled shadow or in direct sunlight. So capacity to see albedo (or pseudo albedo) is just one of the capacities for intentional perception that an achromat has. This is a mysterious capacity to be sure, but it is not a problem particular to albedo perception. Importantly, it is a capacity, like the perception of non-linear motion, that can be absent in a perceiver without rendering that subject *blind*, devoid of *all* visual phenomenology, here ‘black and white’ visual phenomenology. We can imagine a person who has a deficit in motion perception, who sees that a ball has moved from here to there without seeing the ball *move* and indeed such people exist, albeit rarely. It is no less imaginable that a person might suffer a deficit in albedo perception. He or she would see a scene, filled with objects of a certain shape in specific positions of various kinds, but without being able to make judgments of absolute albedo. Similarly a trichromat might see everything that a trichromat can see at night and yet *not* see the surfaces as light or dark. Yes, I made it from the bed to the bath without turning on the

light; no, I did not see whether the hotel carpet was darker than my robe. If this is possible, then seeing albedo is not the basis of seeing 'in black and white'. An achromat or the trichromat at night still *sees* the world, but need not have accurate albedo perception per se.

We end up, then, just where we might have predicted (had we donned our Sellars' hats) in the first place. In retrospect, Nordby's suggestion is quite odd. It is the view that a trichromat who looks at an achromatic photograph (one without chromatic colours) has the same visual experience as an achromat who views a world of many colours. But without prior bias—unless one assumes that our neural processes follow the divide of external images—why would anyone believe that this suggestion *must* be right?

## ACKNOWLEDGEMENTS

I would like thank the James S. McDonnell Foundation through for the Foundation's generous support through a James S. McDonnell Centennial Fellowship, without which this research would not have been possible. I would also like to thank Martin Hahn for reading and commenting on many drafts of this paper; Pete Mandik for his comments and helping to make this paper intelligible to the New York crowd, and; for for their support and time, past and present graduate students, Jason Leardi, Lyle Crawford, Simon Pollon, Emma Esmaili, Nicole Pernat, Gerry Viera, and Robert Foley. Thanks to Richard Brown for organizing and facilitating the Consciousness Online conferences.

## Works Cited

Works cited should go here (this page is optional). See the Works Cited Format document for information on formatting citations.

- Adelson, E. (1993). "Perceptual organization and the judgment of brightness." *Science (New York, NY)*.
- Adelson, E. H. (2000). Lightness Perception and Lightness Illusions. *The New Cognitive Neurosciences*. M. Gazzaniga. Cambridge, MA, MIT Press: 339-351.
- Anderson, B. L. and B.-G. Khang (2010). "The role of scission in the perception of color and opacity." *J Vis* **10**(5): 26.
- Anderson, B. L. and J. Winawer (2005). "Image segmentation and lightness perception." *Nature* **434**(7029): 79-83.
- Anderson, B. L. and J. Winawer (2008). "Layered image representations and the computation of surface lightness." *Journal of vision* **8**(7): 18-18.
- Arend, L. E. and B. Spehar (1993). "Lightness, brightness, and brightness contrast: 1. Illuminance variation." *Percept Psychophys* **54**(4): 446-456.
- Baker, C. L., J. C. Boulton, et al. (1998). "A nonlinear chromatic motion mechanism." *Vision Research* **38**(2): 291-302.
- Buck, S. L., L. P. Thomas, et al. (2006). "Do rods influence the hue of foveal stimuli?" *Vis Neurosci* **23**(3-4): 519-523.
- Cao, D., J. Pokorny, et al. (2011). "Isolated mesopic rod and cone electroretinograms realized with a four-primary method." *Documenta ophthalmologica Advances in ophthalmology* **123**(1): 29-41.
- Cao, D., J. Pokorny, et al. (2008). "Rod contributions to color perception: linear with rod contrast." *Vision Research* **48**(26): 2586-2592.
- Cohen, J. (2004). "Color Properties and Color Ascriptions: A Relationalist Manifesto." *Philosophical Review* **113**(4): 451-506.
- Corney, D. and R. B. Lotto (2007). "What are lightness illusions and why do we see them?" *PLoS Comput Biol* **3**(9): 1790-1800.
- Cropper, S. J., K. T. Mullen, et al. (1996). "Motion coherence across different chromatic axes." *Vision Research* **36**(16): 2475-2488.
- Field, G. D., M. Greschner, et al. (2009). "High-sensitivity rod photoreceptor input to the blue-yellow color opponent pathway in macaque retina." *Nature Neuroscience* **12**(9): 1159-1164.
- Garcia-Suarez, L. and K. T. Mullen (2010). "Global motion processing in human color vision: a deficit for second-order stimuli." *J Vis* **10**(14): 20.
- Gegenfurtner, K. and D. C. Kiper (1992). "Contrast detection in luminance and chromatic noise." *Journal of the Optical Society of America* **9**(11): 1880-1888.
- Gegenfurtner, K. R. (2003). "Cortical mechanisms of colour vision." *Nat Rev Neurosci* **4**(7): 563-572.
- Gheorghiu, E. and F. A. Kingdom (2007). "Chromatic tuning of contour-shape mechanisms revealed through the shape-frequency and shape-amplitude after-effects." *Vision Res* **47**(14): 1935-1949.
- Greenlee, M. W., S. Magnussen, et al. (1988). "Spatial vision of the achromat: spatial frequency and orientation-specific adaptation." *J Physiol (Lond)* **395**: 661-678.
- Hansen, T. and K. R. Gegenfurtner (2009). "Independence of color and luminance edges in natural scenes." *Vis. Neurosci.* **26**(1): 35-49.
- Hess, R. F. and K. Nordby (1986). "Spatial and temporal limits of vision in the achromat." *J Physiol (Lond)* **371**: 365-385.
- Jacobs, G. H. (2008). "Primate color vision: A comparative perspective." *Vis. Neurosci.* **25**(5-6): 619.
- Jacobs, G. H. (2009). "Evolution of colour vision in mammals." *Philos Trans R Soc Lond, B, Biol Sci* **364**(1531): 2957-2967.
- Jacobs, G. H. and M. P. Rowe (2004). "Evolution of vertebrate colour vision." *Clinical & experimental optometry : journal of the Australian Optometrical Association* **87**(4-5): 206-216.
- Kingdom, F. A. A. (2003). "Color brings relief to human vision." *Nat Neurosci* **6**(6): 641-644.
- Kingdom, F. A. A. (2005). "Chromatic properties of the colour-shading effect." *Vision Research* **45**(11): 1425-1437.

- Kingdom, F. A. A. (2008). "Perceiving light versus material." *Vision Res* **48**(20): 2090-2105.
- Kingdom, F. A. A. (2011). "Lightness, brightness and transparency: a quarter century of new ideas, captivating demonstrations and unrelenting controversy." *Vision Research* **51**(7): 652-673.
- Kingdom, F. A. A. and R. Kasrai (2006). "Colour unmask dark targets in complex displays." *Vision Res* **46**(6-7): 814-822.
- Kingdom, F. A. A., K. Wong, et al. (2006). "Colour contrast influences perceived shape in combined shading and texture patterns." *Spatial vision* **19**(2-4): 147-159.
- Knill, D. C. and D. Kersten (1991). "Apparent surface curvature affects lightness perception." *Nature* **351**(6323): 228-230.
- Li, W., S. Chen, et al. (2010). "A fast rod photoreceptor signaling pathway in the mammalian retina." *Nat Neurosci* **13**(4): 414-416.
- Michna, M. L., T. Yoshizawa, et al. (2007). "S-cone contributions to linear and non-linear motion processing." *Vision Research* **47**(8): 1042-1054.
- Mullen, K. T., W. H. Beaudot, et al. (2000). "Contour integration in color vision: a common process for the blue-yellow, red-green and luminance mechanisms?" *Vision Res* **40**(6): 639-655.
- Mullen, K. T. and W. H. A. Beaudot (2002). "Comparison of color and luminance vision on a global shape discrimination task." *Vision Res* **42**(5): 565-575.
- Nordby, K. (1996). "Vision in a Complete Achromat: A Personal Account." from <http://consc.net/misc/achromat.html>.
- Olkkonen, M., T. Hansen, et al. (2008). "Color appearance of familiar objects: effects of object shape, texture, and illumination changes." *J Vis* **8**(5): 13.11-16.
- Pang, J.-J., F. Gao, et al. (2010). "Direct rod input to cone BCs and direct cone input to rod BCs challenge the traditional view of mammalian BC circuitry." *Proc Natl Acad Sci USA* **107**(1): 395-400.
- Poirier, F. J. A. M. (2009). "The Anderson-Winawer illusion: it's not occlusion." *Atten Percept Psychophys* **71**(6): 1353-1359.
- Sellars, W. (1956). *Empiricism and the Philosophy of Mind*. Minneapolis, MN, University of Minnesota Press. **I**: 253-359.
- Sellars, W. (1962). *Philosophy and the Scientific Image of Man*. *Frontiers in Science and Philosophy*. R. Colodny. Pittsburgh, PA, Pittsburgh Press: 35-78.
- Skottun, B., K. Nordby, et al. (1982). "Temporal summation in a rod monochromat." *Vision Research* **22**(4): 491-493.
- Spillmann, L., J. Hardy, et al. (2010). "Brightness enhancement seen through a tube." *Perception* **39**(11): 1504-1513.
- Stabell, B. (1967). "Rods as color receptors in scotopic vision." *Scand J Psychol* **8**(2): 132-138.
- Stabell, U. (1967). "Rods as color receptors in photopic vision." *Scand J Psychol* **8**(2): 139-144.
- Stabell, U. and B. Stabell (1965). "Rods as color receptors." *Scand J Psychol* **6**(3): 195-200.