Retinex-like computations in human lightness perception and their possible realization in visual cortex

Michael E. Rudd; University of Washington; Seattle, WA/USA

Abstract

Land and McCann's original Retinex theory [1] described how the brain might achieve color constancy by spatially integrating the outputs of edge detector neurons in visual cortex (i.e., Hubel and Wiesel cells). Given a collection of reflective surfaces, separated by hard edges (a Mondrian stimulus) and viewed under uniform illumination, Retinex first computes luminance ratios at the borders between surfaces, then multiplies these ratios along image paths to compute the relative ratios of noncontiguous surfaces. This multiplication is equivalent to summing steps in log luminance. Here I review results from the human lightness literature supporting the key Retinex assumption that biological lightness computations involve a spatial integration of steps in log luminance. However, to explain perceptual data, the original Retinex algorithm must be supplemented with additional perceptual principles that together determine the weights given to particular image edges. These include: distance-dependent edge weighting, different weights for incremental and decremental luminance steps, contrast gain acting between edges, top-down control of edge weights, and computations in object-centered coordinates. I outline a theory, informed by recent findings from neurophysiology, of how these computations might be carried out by neural circuits in the ventral stream of visual cortex.

Edge integration in Retinex theory

Figure 1 shows a collection of achromatic papers of different sizes and shapes, separated by hard edges. This type of stimulus is sometimes referred to as a Land Mondrian pattern [2], after the painter Piet Mondrian who created paintings comprised of rectangular patches. When the magnitude of the uniform illumination lighting such a pattern is varied, the perceived paper reflectances (i.e., their *lightnesses*) remain remarkably stable, even though the luminance of each paper scales in proportion to the illumination level. Clearly, lightness ≠ luminance.

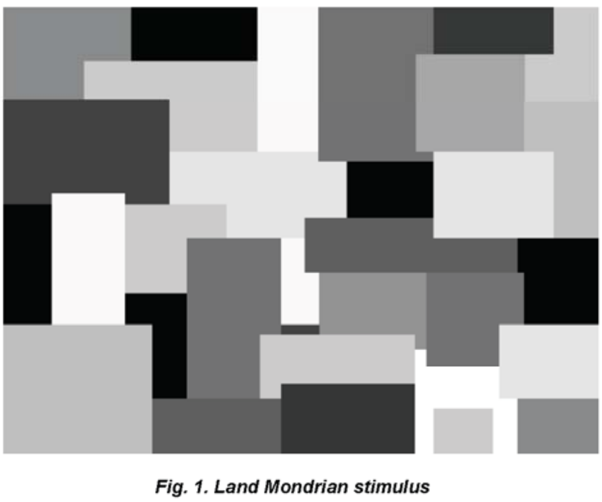

Fig. 1. Land Mondrian stimulus

In their original Retinex theory, Land and McCann [1] proposed a model of how the brain—or a technological imaging device—might achieve color constancy under the challenge of varying illumination level and/or wavelength composition. In this paper, I constrain my discussion to the problem of lightness (i.e., achromatic color) constancy under varying illumination level. Retinex achieves lightness constancy by first computing the local luminance ratios L_{n+1}/L_n at the borders between contiguous surfaces—information that might be provided by edge detector neurons in visual cortex—then multiplies ratios along paths through the image to compute the luminance ratios of any two arbitrary pairs of surfaces *i* and *j*:

$$\prod_{n-i}^{j-1} \frac{L_{n+1}}{L_n} \quad . \qquad (1)$$

Multiplying luminance ratios along an image path is mathematically equivalent to spatially summing local steps in log luminance at borders. Thus, an alternative method for achieving lightness constancy is to log transform the image, then sum steps in log luminance:

$$\sum_{n-1}^{j-1} \left(\log L_{n+1} - \log L_n \right) \quad . \qquad (2)$$

In what follows, I will refer to an operation that spatially sum steps in log luminance across space for the purpose of computing lightness as *edge integration*. Edge integration suffices to determine the lightness scale for the whole image up to an unknown constant. To compute an absolute lightness scale—that is, the particular lightnesses that we actually *perceive*—requires an additional operation that serves to fix the value of this constant. This second operation is sometimes referred to as *lightness anchoring*. Various anchoring rules have been proposed in the literature to account for perceived lightness. These include: the rule that the highest luminance always appears to be white (say, a surface having a 90% reflectance [1,3]), or that the average image luminance is perceived as gray [4,5]. I have suggested that the rule that is actually employed by our visual system may be that the highest *lightness* is always seen as white [6-10]. This highest lightness anchoring rule can be distinguished from the highest luminance rule because edges are typically summed sub-additively [6-10]. The question of anchoring is further complicated by the fact that surfaces sometimes appear self-luminous and thus are not assigned a shade of gray [11]. The problem of self-luminosity, while important for models of lightness computation, will not be further addressed here.

Violations of lightness constancy and the breakdown of edge integration

The tendency for perceived surface reflectance to remain stable under changes in illumination level is sometimes referred to as *Type I lightness constancy* to distinguish this type of constancy from *Type II lightness constancy*, or the stability of surface lightness under a change in the spatial context in which the surface is viewed [11]. Type I constancy holds to a good degree of approximation in human vision, while Type II constancy does not.

A classic example of a *failure* of Type II constancy is *simultaneous contrast*, in which a mid-luminance target patch appears lighter when viewed against a dark background than an identical target patch viewed against a light background [12]. As the name implies, this phenomena is often explained on the basis of local contrast alone (or the local luminance ratio at the target/background border). However, a number of results in the lightness literature demonstrate that it is not only the local contrast at the target border that contributes to the computation of lightness, but rather that lightness is computed from a sum of steps in log luminance across space—as in Retinex—in which the local ratio at the target border simply makes a stronger contribution to the sum than do borders that are more distant from the target [6-15].

An illustration of this idea—which I will refer to here as *distance-dependent edge integration*—is presented in Figure 2. The two disks and in the figure have the same luminance, as do the annuli that surround them. Thus, the local luminance ratio at the disk/annulus border is the same on the two sides of the figure, but the disk and ring on the left *both* appear lighter than on the right.

Fig. 2. Demo of perceptual edge integration in lightness

We have already seen that lightness ≠ luminance. Figure 2 suffices to demonstrate that lightness ≠ local contrast. Nevertheless, the difference in the perceived lightness of the disks on the two sides of Figure 2 *can* be understood in terms of an edge *integration* process in which lightness is computed from a spatial sum of steps in luminance across space, but in which the luminance step at the local edge between the disk and its surround annulus makes a stronger contribution to the disk lightness then does the step at the more distant annulus/background edge.

Figure 3 demonstrates more directly that distance plays a critical role in determining the contribution of an edge step to lightness. Again, the disks and annuli on the two sides of the display have the same luminances, but the disk on the right appears lighter. And again, the percept evoked by Figure 3 can also be understood in terms of an edge integration process in which luminance steps are perceptually integrated to compute the disk lightness, but the step at the local disk/annulus border makes a stronger contribution to the disk lightness than does the more distant annulus/background border. Furthermore, the step at the annulus/background border makes a stronger contribution to the disk lightness on the right side of the display than on the left (in both cases, signaling that surfaces on the inside of the annulus/background border are lighter than surfaces on the outside) because it is closer to the target disk.

The idea that lightness depends on a long-range sum of local contrasts, in which the weights given to contrast elements decline with distance was first proposed by Reid and Shapley [13,17], who developed and tested an edge integration model based on weighted sums of Michelson contrast. Later, Rudd and Zemach [6,7] developed a closely related quantitative edge integration model in which lightness is computed from sums of directed steps in log luminance, rather than from sums of Michelson contrast. They showed that the model based on steps in log luminance produced a better fit to perceptual data. They then went on to show that the strength of induction from an annular surround depends not only the relative distances of the two annulus borders, but also on the contrast polarities of the two borders [6,7,15]; see also [8-10].

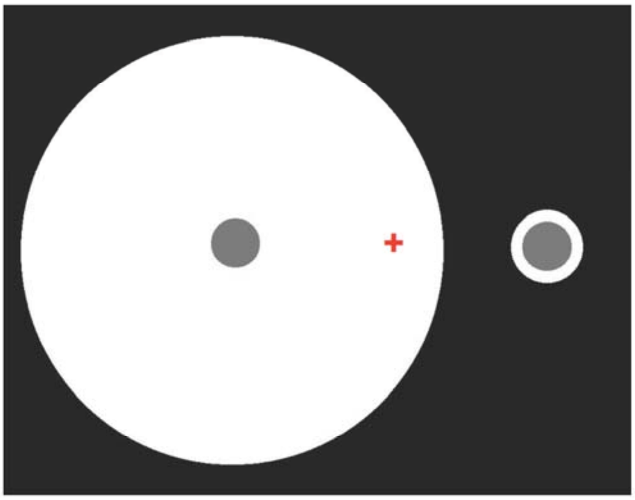

Fig. 3: Distance-dependent edge integration demo

Rudd [9] subsequently showed that all of the perceptual effects discussed thus far can be explained to a high degree of quantitative precision (the percentages of variance in lightness matches explained often exceed 99%) with an edge integration model in which the lightness Φ_j of a target disk (designated here as Surface "*j*"), that is embedded in a series of concentric surround surfaces, labeled 1, 2, 3,…, *j*-1, is computed from a weight sum of steps in log luminance evaluated along a path from the background (here designated Surface 1) to the target, according to the edge integration algorithm

$$\Phi_j = \sum_{n=1}^{j-1} w_{n+,n} \left(\log L_{n+1} - \log L_n\right) , \qquad (3)$$

where the weight given to a border between any two surfaces *n* and *n+1*, in computing the disk lightness, depends on two *independent* factors. The first is the distance of the edge from the target. The second is the edge contrast polarity, defined in the direction of the target; that is, whether the edge step entails an increase, or a decrease, in luminance, along a directed path from the background and the target disk.

This two-factor model of lightness computation can be viewed as a modified Retinex model in which both edge locality and edge contrast polarity play a role in determining the contributions that individual edges make to the lightness of target surfaces embedded in a complex scene. But it is not yet a viable general model of lightness computation, even in 2D images [9,10]. Furthermore, while distance and edge polarity may seem at first glance to be low-level features, the situation appears to be more complex than that. The fact that the edge integration computation proceeds in only one direction—from the background to the target in the case of disk-annulus stimuli—and that edge contrast polarity is defined relative the this path, both suggest that the computations described by the two-factor model are not 'low-level" in the sense of a computation that depends only on local image features, but rather that it takes place in target-centered coordinates [10]. This conclusion is further reinforced by need for a third modification of the original Retinex model that I have found necessary to make in order to explain perceptual data—a modification that also depends on target-centered coordinates. Before describing this third modification, I first need to review the literature on lightness assimilation.

Combined effects of edge-based contrast and "assimilation" on lightness

In the relatively simple case of a disk surrounded by a single annulus, the lightness computation algorithm specified by Eq. (3) takes the form

$$\begin{aligned}\Phi_D &= w_{D,A}\left(\log L_D - \log L_A\right) + w_{A,B}\left(\log L_A - \log L_B\right)\\ &= w_{D,A}\log L_D - w_{A,B}\log L_B - \left(w_{D,A} - w_{A,B}\right)\log L_A\end{aligned} \qquad (4)$$

where the symbol Φ_D represents the disk lightness; L_D, L_A, and L_B are the luminances of the disk, annulus and background; and $w_{D,A}$ and $wA.B$ are the perceptual weights given to the disk/annulus and annulus/background edges in computing the disk lightness.

It is worth pausing to unpack some of the implications of Eq. (4). The top line of the equation emphasizes the fact that the disk lightness depends on two separate components: one corresponding to the weighted luminance step (in log units), $w_{D,A}(\log L_D\text{-}\log L_A)$, at the disk/annulus edge, and the other corresponding to the weighted luminance step (in log units), $w_{A,B}(\log L_A\text{-}\log L_B)$, at the annulus/background edge. When the inner and outer edges of the annulus have different contrast polarities, these two terms will have opposite signs. The two disk-annulus pairs shown in Figure 3 provide examples of such a situation. Here the step in log luminance, $\log L_D\text{-}\log L_A$, at the disk/annulus edge, measured along a path pointed inward towards the disk whose lightness is being computed, is *negative* because the disk luminance less than the annulus luminance; whereas, the step in log luminance at the annulus/background edge, $\log L_A\text{-}\log L_B$, is *positive* because the annulus luminance is greater than the background luminance. The two edges thus have opposing effects on the disk lightness. The inner edge induces darkness in the disk, while the outer edge induces lightness.

As mentioned above, we can think of the presence of the luminance step at the inner edge as informing the visual system that surfaces located on the inside of this edge are darker than surfaces on the outside of this edge; while the presence of the luminance step at the outer edge informs the visual system that surface located on the inside of the outer edge are lighter than surfaces located on the outside of the outer edge, information that is applies perceptually to both the annulus and the disk.

Possibly because of the fact that the inner and outer edges of the annulus can have opposite effects on the disk lightness, Reid and Shapley [13,14] referred to the outer edge as exerting an "assimilation" effect on the disk lightness. However, their use of the term "assimilation" to denote the contribution of the outer ring edge to the computation of the disk lightness may lead to confusion because the term *assimilation* has another, more generally accepted, definition in the context of lightness perception. Classically [16,17], the term assimilation was used to denote cases in which bordering a region with another lighter region has the effect of lightening the first region. This classical assimilation sometimes also goes by the name of reversed, or paradoxical, contrast because it has an effect on target lightness that is in the direction opposite to contrast. In any case, the edge integration algorithm represented by Eq. (4) holds regardless of the contrast polarities of the inner and outer annulus edges, including in cases where the steps in log luminance at the inner and outer edges of the annulus are both positive or both negative and thus the two edges do not have opposite influences on the disk lightness.

When the inner and outer annulus edges *do* have opposite contrast polarities, as in Figure 3, the polarity that "wins out" in determining the effect that the annulus luminance has on the disk lightness depends on which edge weight—the weight associated with the inner edge or the weight associated with the outer edge—is largest. To this point, I have assumed the validity of the two-factor model that says that edge weights decline with distance from the target and that the weights associated with positive contrast polarity edges (i.e., increasing luminance steps in the direction of the target) are only about 1/3 as large as the weights associated with negative contrast polarity edges. It follows that, for the case of the stimuli in Figure 3, $w_{A,B} < w_{D,A}$. So, the second line of Eq. (4) tells us that increasing the annulus luminance L_A will have the effect of decreasing the disk lightness for this particular combination of inner and outer edge contrast polarities (which is what happens [6]). Thus, local contrast wins out over the countering influence of the outer (positive) edge step and the net influence of the surround annulus on disk lightness is one of contrast.

It is worth noting here that the original Retinex model of Land and McCann will produce no such net contrast effect in this case because their Retinex theory assumed that all edge weights are equal and thus that $w_{A,B} = w_{D,A}$ in Eq. (4). The original Retinex must be modified in order to account for simultaneous contrast.

The need to included contrast gain control to explain lightness assimilation

In the above discussion, I referred to distinct definitions of the term "assimilation." It is worth giving these different types of assimilation labels in order to keep things clear. The first type of assimilation is the classical type in which neighboring or surrounding a target region with a region of higher luminance lightens the target region [16,17]. I will refer to this classical type of assimilation as Type I assimilation. Since this is the standard usage in the lightness community for the term assimilation, we should reserve the term "assimilation," when unqualified, to refer to Type I assimilation.

The second type of assimilation, which I will call Type II assimilation, or Reid-Shapley assimilation, refers to the independent effect on target lightness of the outer annulus border.

A third—and also theoretically import—type of assimilation was discovered in the course of my own experiments on lightness induction using disk-annulus stimuli [6-10,15,18-20]. I will refer to this third type of assimilation as Type III assimilation.

According to edge integration theory, the magnitude to the induction effect produced by a log-unit change in annulus luminance on the disk lightness is given by the following partial differential equation

$$\frac{\partial \Phi_D}{\partial \log L_A} = -\left(w_{D,A} - w_{A,B}\right) \ . \qquad (5)$$

Eq. (5) implies that a plot of disk lightness versus annulus luminance should be a straight line whose slope depends on the weights associated with the inner and outer annulus edges.

This model prediction holds approximately for the data from some lightness matching experiments conducted with disk-annulus displays, but significant violations of the model are also often observed [6-10,15,18-20]. A case in point is illustrated in Figure 5. Here, the particular stimulus was a disk-annulus pair whose inner and outer annulus edges both had negative contrast polarities (see [8] for details). However, quantitative analyses have shown that similar violations of the straight-line prediction of Eq. (5) occur for with all disk-annulus displays. These violations always take the form of curvature in lightness plot that is not predicted by Eq. (5), but that can be alternatively accounted for by a model in which lightness varies as a parabolic (second-order polynomial) function of annulus luminance. When the linear model (Eq. (5)) does approximately hold, it is only because the second-order term in the parabolic model is negligible (see above references, as well as the parabolic regression models of the data in Figure 5).

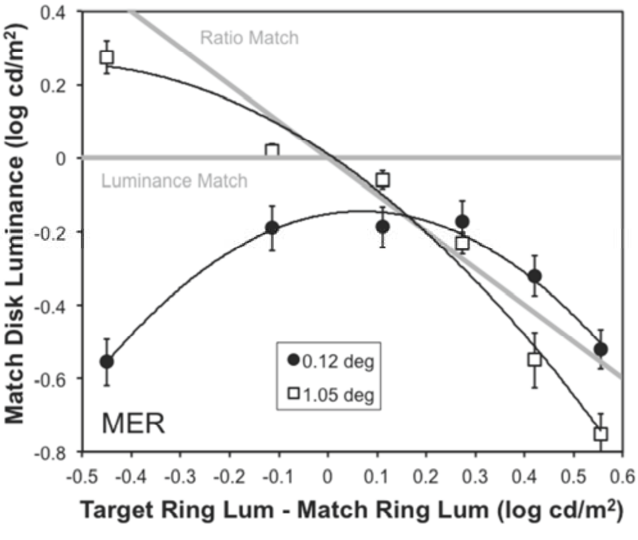

Fig. 5. Examples of parabolic lightness functions

The curvilinear trend in lightness functions, such as the ones shown in Figure 5 and others studied in my lab [8, 15], tends to be most pronounced when the annulus width is small (e.g., the 0.12 deg annulus plot in Figure 5). In such cases, there may be a significant range of annulus luminances for which increases in the annulus luminance produces increases, rather than decreases, in the disk lightness. It is this situation that I am here calling Type III assimilation because increasing the ring luminance has the effect of lightening the target, an effect that is in the direction opposite to contrast. Interestingly, Type I assimilation—classical assimilation also tends to be produced when target regions are neighbored by narrow bright strips or borders. So the two types of assimilation are likely related. As Figure 5 illustrates, a stimulus that produces Type III assimilation over one range of annulus luminance may produce contrast over another range. Thus, the two phenomena—contrast and assimilation—appear to be part of a larger overarching pattern of lightness induction.

Intuition suggests that the fact that the slope of the disk lightness plot can be positive over one range of annulus luminance, but negative over another range might mean that the perceptual weights associated with edges actually depend on the ring luminance. Furthermore, the fact that the plot curvature is more pronounced when the annulus is narrow suggests the possibility that the presence of another, nearby, edge influences the perceptual weight given to an edge.

It is known from cortical neurophysiology that the neural gain of edge detector neurons in early visual cortex can be influenced by the activation of other nearly edge detector neurons. This effect is known as "contrast gain control." Rudd [8] showed that the parabolic effects seen in all available data from lightness experiments employing disk-annulus stimuli can be explained by an edge integration model supplemented with a contrast gain control mechanism. In this supplemented edge integration model, the luminance ratio of the outer annulus edge modulates the neural gain (and hence the perceptual weight) applied to the inner edge, and vice versa. A detailed analysis of the parabolic lightness plots from these experiments showed that increasing the (absolute) size of the luminance step associated with the outer edge tends to *decrease* the perceptual weight assigned to the inner edge in the computation of disk lightness; whereas, increasing the (absolute) size of the luminance step associated with the outer edge tends to *increase* the perceptual weight assigned to the inner edge [8].

Thus, to explain the lightness percepts with a modified Retinex theory (i.e., an edge integration model), it must be assumed that the sign of the contrast gain control knows about "inner" and "outer" edges and operates in an object-centered framework [8-10]. This finding reinforces the conclusion that the lightness computations studied here are carried out at a mid-level stage of visual analysis, which implies a cortical locus for neural lightness computation.

The addition of the contrast gain control to the edge integration model means that the edge weights in Eq. (4) are now modeled by the considerably more complex expressions

$$\begin{aligned} w_{D,A} &= \omega(x_{D,A})g(p_{D,A})[1 - f(d)g(p_{A,B})\left|\log L_A - \log L_B\right|]^+; \\ w_{A,B} &= \omega(x_{A,B})g(p_{A,B})[1 + h(d)g(p_{D,A})\left|\log L_D - \log L_A\right|]^+, \end{aligned} \qquad (6)$$

where the function $\omega(x)$ describes the spatial fallout in the contribution of an edge to the disk lightness; g is the neural gain factor applied to an edge, which according to the two-factor model depends on the edge contrast polarity and is +1 for edges whose dark sides point towards the disk and +1/3 for edges whose light sides point towards the disk; $f(d)$ and $h(d)$ are functions that describe the spatial falloff in the strengths of contrast gain control acting between edges; d is the annulus width (distance between the inner and outer annulus edges); and the mathematical operator $[Z]^+$ denotes half-wave rectification, or $max[0,Z]$.

The slope (Eq. (5)) of the lightness plot now depends on the relative balance of the two weights described by Eq. (6) and can be either negative or positive, depending on the annulus luminance. The plot will have an inflection point when the luminance is such that the two edge weights are equal.

Neural interpretation of the computational model

Edge detection occurs in the brain as early as the simple cells of primary visual cortex (also known as area V1) [21-23]. Simple cells are often modeled as rectified oriented 2D linear spatial filters, having Gabor-like receptive fields (i.e., sinusoids modulated by a Gaussian envelope) [24,25]. The receptive fields of actual simple cells span a range of sinusoidal phases with respect to the Gaussian envelope, including the canonical even- and odd-symmetric (cosine and sine) phases. For the purpose of discussing edge detection, it suffices to consider only odd-symmetric receptive fields in what follows.

Simple cell receptive fields are constructed from the outputs of neurons at earlier visual processing stages (retinal gain cells, LGN cells, neurons in the input layer 4 to V1), whose receptive fields are circularly symmetric and thus are not selective for edges of a particular orientation or contrast polarity. These circularly symmetric receptive fields come in two types, ON-center and OFF-center, which selectively respond to incremental and decremental luminance [23]. Recent neurophysiological evidence indicates that OFF-center neurons respond more strongly to differential luminance than do ON-center neurons, which implies that the inherent gains of OFF-center cells is larger than that of ON-center cells [26]. This allows for the possibility that simple cell receptive fields could be constructed to respond to edges of different contrast polarities with different inherent gains [8,9].

I am currently developing a computer implementation of this idea, but detailed simulations are beyond the scope of this short paper. Suffice it to say that mid-level cortical computations that spatially integrate the outputs of V1 simple cells responding to edges different contrast polarities could, in principle, selectively integrate the outputs of neurons whose firing rates reflect the different inherent gains of the ON- and OFF-cells from which their receptive fields are constructed, depending on whether an ON- or OFF-center receptive field that drives the simple cell to fire lies on the side of the edge that points in the direction of the target disk. On this theory, the neural gain applied to the edge in the process of cortical lightness computation inherits the gain of the ON- and OFF- units from which the edge response is constructed [8,9].

As mentioned above, evidence exists for additional modulation of the neural gains of simple cells by the responses of other, nearby, simple cells. Eqs. (6) can thus be interpreted in terms of a mechanism by which the neural response to an edge in cortical Neuron 1 depends on that neuron's own inherent gain, but is also further modulated by the response of a nearby Neuron 2, in a way that depends on both the cortical distance between Neurons 1 and 2, and the inherent gain of Neuron 2. However, whether the edge detector neurons in the computational lightness model should be identified with the simple cells in area V1 is presently unclear, because the sign of the contrast gain control, as mentioned above, depends on the direction of the contrast gain control relative to the center of the surface whose lightness is being computed, which is not a low-level local image feature. This may implicate the involvement of higher stage of cortical processing, such as area V2, where edge detector neurons are also known to exist, but where neural responses depend on mid-level perceptual constructs, such as image segmentation and perceptual organization [8,9].

In any case, the computational model requires the existence an additional cortical stage that integrates outputs of the neural populations in early visual cortex that respond to the inner and outer edges of the ring, after the neural gains applied to these edges have been subjected to the various gain modulations described above. I have suggested that this higher, edge integration, stage may be identified with separate 'lightness' and 'darkness' neurons in area V4 and that the lightness that we perceive depends on a comparison of the outputs of these lightness and darkness neurons at a subsequent cortical stage, such as cortical area TE or TEO [8-10].

It remains an open question for the theory whether the logarithmic transformation of luminance required by the model occurs before or after the spatial filtering stage of the simple cell receptive fields. Since the formation of simple cell receptive fields actually begins with the center-surround organization of the ON- and OFF- bipolar cells in the retina, a logarithmic transformation that came before receptive field spatial filtering would have to occur in the photoreceptors. This would imply that the simple cell receptive fields, instead of acting as spatial differentiators of *luminance* per se, would act as spatial differentiators of *log luminance*, which is an operation equivalent to taking the logs of ratios, as required by the edge integration model. A visual system that log transformed image luminance before applying spatial filtering would be capable of computing log of luminance ratios at different spatial scales and orientations, by making use of different scales and orientation of the Gabor filters that model that simple cell receptive fields [8-10,18]. However, physiological evidence seems to refute the idea that either the rod or cone photoreceptors perform a logarithmic transformation of input intensity. Furthermore, the same computational function could be achieved, as least in cone vision, by a log transformation of the simple cell output in combination with the Weber ratio encoding property of the cone photoreceptors, which are ignored in the Gabor model of the simple cell, but are nevertheless well documented.

Generalizing the lightness model to explain a novel luminance gradient illusion

In a recent paper [27], I applied this neural theory to the explain a novel—a otherwise puzzling—luminance gradient illusion: the *Phantom Illusion* (Figure 6). To produce the illusion, either an incremental or a decremental target square is surrounded by a shallow linear luminance gradient whose luminance either decreases in the direction of the target (for incremental targets) or increases in the direction of the target (for decremental targets). For sufficiently wide gradients, incremental and decrements targets appear veridically as increments or decrements. For sufficiently narrow gradients, however, increments appear as decrements and vice versa. The latter situation can be thought of as yet another example of assimilation, in which the relatively steep gradient induces a lightness effect in the target that overwhelms the local effect of edge contrast. Nevertheless, this illusion can be explained by the same edge integration model that explains Type III assimilation in disk-annulus experiments.

To see this, note that cortical neurons that respond to edges would also be expected to respond to gradients. Specifically, the half-wave rectified output of an oriented Gabor filter having an odd-symmetric kernel will produce a positive response both to edges having a particular contrast polarity and to gradients whose luminance ramps are oriented in the direction for which the filter is selective. Filters responding to an edge will produce a output that is proportional to the edge step. Filters responding to a gradient will produce an output that is proportional to the gradient slope (see Ref. [27] for further details).

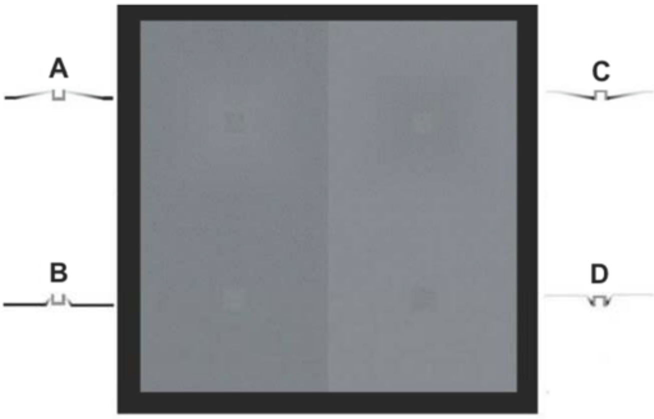

Fig. 6. The phantom illusion. Squares A and B are both luminance decrements; while C and D are luminance increments. Surrounding either with a narrow luminance ramp reverses the apparent target polarity.

To model the Phantom Illusion, I constructed a 1D computer simulation of the response to the gradient-embedded targets of the lightness computation model corresponding to Eqs. (4) and (6), where the functions $\omega(x)$, $f(d)$, and $g(d)$ were all assumed to be exponentially decaying functions of distance. In the cases of $f(d)$ and $g(d)$, these functions were generalized to describe distances between neurons responding to both edges and gradients. The simulation results (Figure 7) correctly predicted that sufficiently shallow gradients produce target lightness percepts dominated by local edge contrast, while sufficiently steep gradients produce lightness percepts dominated by the sum of the neural responses within the gradient. The critical insight of the model is that the target percept depends not on local contrast per se, but rather on a spatial integration of the outputs of local 'edge' (or gradient) detectors that respond to spatial variations in illumination at some orientation (e.g., 'edge' integration). That is, the lightness that we perceive is neither local luminance, nor local contrast, but rather the output of a spatially-extended lightness computation based on the long-range spatial integration of local oriented contrasts.

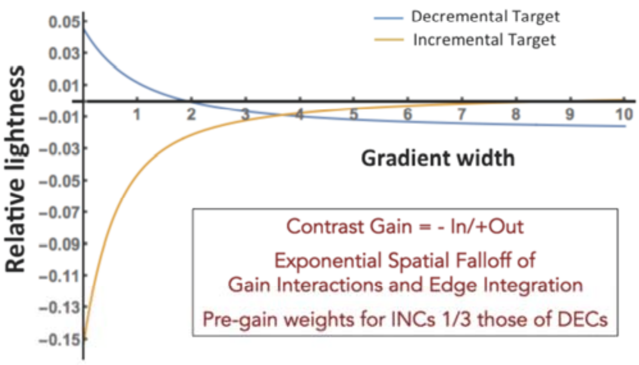

Fig. 7. Computer simulation of the Phantom Illusion based on the lightness model. When surrounded by wide luminance gradients, incremental and decremental targets are seen veridically as increments and decrement. Decreasing the gradient width produces a reversal of the apparent target polarity. See Ref. [27] for further details.

Edge classification and the influence of attention on lightness computation

A final point is worth mentioning here because it potentially speaks to cortical processing level at which these lightness computations occur. This point is that the observer's *interpretation* of local contrast elements within the scene can influence the output of their perceptual lightness computations *even when the physical stimulus remains constant*. From a functional standpoint, the edge integration computations carried out by any Retinex-like model should ideally only spatially integrate luminance steps in the image that result from changes in surface reflectance in the external world. Spatially integrating steps in surface reflectance servers to establish a global lightness scale that applies to all of the surfaces within the scene. On the other hand, any tendency that the visual system might have to integrate steps in luminance in the image that are actually due to spatial variation in illumination, rather than reflectance, will tend to produce a false lightness scale. To put it another way, one primary goal of any lightness constancy algorithm should be to 'discount the illuminant,' as Helmholtz [28] famously suggested that we do.

The extent to which our visual system only integrates steps in reflectance, and not illumination, is beyond the scope of this paper. But a study by Rudd [10] demonstrates that subjects are at least to some extent able to control the edge integration process by selectively integrating only edges that they conceive of as reflectance edges. The subjects in that study were instructed to interpret the same hard outer annulus edge in a disk-annulus display as being due to either a change in surface reflectance or illumination. Subjects who interpreted the edge as resulting from illumination change were able to suppress the edge from entering into their perceptual computations of lightness, as judged from fits of the edge integration model to their data. Furthermore, this suppression resulted in the elimination of the effects of contrast gain control between edges. The results have a number of important implications for the neural and perceptual interpretations of the computational model. First, they suggest that attention (or intention) can play a role in modulating the neural response to edges in early visual cortex. Second, they reinforce that conclusion that the cortical processes that instantiate the model computations are concerned with cognitive interpretation, rather than simply low-level image processing.

My current thinking about how this cognitive influence might be realized in the brain in schematized in Figure 8, which illustrates the model's processing of a double-increment disk-annulus display (i.e., one whose disk luminance exceeds that of the ring). According to the model, edge detector neurons in cortical areas V1 and V2 selectively encode the step in log luminance at each edge of the display. These neurons then influence each other's neural gains in a way that is described by Eqs. (6) above. In Figure 8, I imagine this as a feedforward process that could, in principle, occur in the feedforward pathway from V1 to V2. But other neural architectures might be proposed to instantiate the same computational model.

Once the contrast gain control has taken place (probably no later than in area V2), a subsequent cortical stage spatially integrates the outputs of the second-stage edge detector neurons whose gains have been influenced both by any low-level gain factors that differ for incremental and decremental edge steps, and by the contrast gain control that takes place between nearby edge detector units. A likely cortical locus for the edge integration operation is area V4, which receives direct input from V2. I place the edge integration stage in V4 because there are known to be distinct classes of neurons in V4 that respond to incremental and decremental luminance (known as 'brightness' and 'darkness' neurons) [29], and these neurons have receptive fields that are sufficiently wide to integrate edge steps over a considerable spatial range. In fact, there are also neurons in V4 that could represent additional axes of color space. Furthermore, cortical damage to V4 results in visual deficits that some neuroscientists have ascribed to a failures of 'Retinex-like' color constancy mechanisms [30-32].

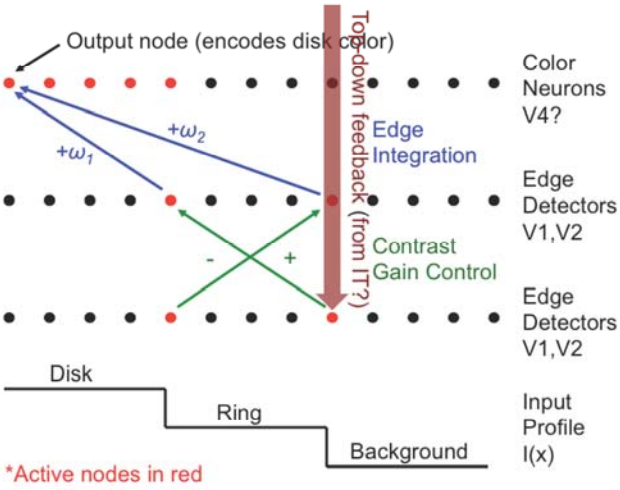

Fig. 8. Current version of the cortical lightness model

To account for the influence of cognitive edge interpretation on the outcome the results of lightness matching experiments conduced with disk-annulus stimuli, I envision a top-down gain control mechanism—most likely originating in IT, or else in pre-frontal cortex, then routed through IT [33]—that reaches all the way down to the earliest stage of cortical processing of the model (i.e., the first edge detection stage in V1 or V2) to modulate the neural gain applied to edges, prior to stage of contrast gain control. It is necessary to assume that the top-down modulation occurs prior to the contrast gain control stage because edge classification has the effect on lightness matching, not only of suppressing the response to an edge that is cognitively classified as being an illumination edge, but also potentially of reversing the sign of the curvature in the lightness plots, turning inverted-U shaped lightness plots, such as the ones shown in Fig. 5, upside down (see Ref [100] for examples and an explanation).

These effects of edge classification and top-down influence are quantitatively modeled by further modifying the equations for perceptual weights applied to edges in the following manner:

$$w_{D,A} = \omega_{D,A}\zeta_{D,A}g\left(p_{D,A}\right)\left[1 - f(d)\zeta_{A,B}g\left(p_{A,B}\right)\left|\log L_A - \log_B\right|\right]^{+};$$
$$w_{A,B} = \omega_{A,B}\zeta_{A,B}g\left(p_{A,B}\right)\left[1 - h(d)\zeta_{D,A}g\left(p_{D,A}\right)\left|\log L_D - \log_A\right|\right]^{+}, \quad (7)$$

where the symbols $\zeta_{D,A}$ and $\zeta_{A,B}$ represent the top-down gain control that selects edges for integration on the basis of subjective edge classification.

To achieve lightness constancy, edges that are classified as reflectance edges should ideally be assigned the top-down gain value $\zeta_{i,j} = 1$; while edges that are classified as illumination edges should be assigned a value $\zeta_{i,j} = 0$. Where this ideal situation actually pertains in human vision or neurophysiology is an important open question for future research.

References

[1] E. Land & J. J. McCann, "The retinex theory of vision," J. Opt. Soc. Am., Vol. 61, 1-11, 1971.

[2] E. H. Land, " The retinex theory of color vision," Sci. Amer., Vol. 237, 108-128, 1977.

[3] A. Gilchrist, Seeing Black and White. New York: Oxford Univ. Press, 2006.

[4] D. B. Judd, "Hue, saturation, and lightness of color surfaces with chromatic illumination. J. Opt, Soc. Am., Vol. 30, 2-32, 1940.

[5] H. Helson, "Adaptation-level as a basis for quantitative theory frames of reference," J. Opt. Soc. Am, Vol. 35, 555-567.

[6] M. E. Rudd & I. K. Zemach, "Quantitative properties of achromatic color induction: An edge integration analysis," Vision Res., Vol 44, 971-981, 2004.

[7] M. E. Rudd & I. K. Zemach, "The highest luminance rule in achromatic color perception: Some counterexamples and an alternative theory," J. Vision, Vol. 5, 983-1003, 2005.

[8] M. E. Rudd, "How attention and contrast gain control interact to regulate lightness contrast and assimilation," J. Vision, Vol. 10, 40, 2010; doi:10.1167/10.14.40.

[9] M. E. Rudd, "Edge integration in achromatic color perception and the lightness-darkness asymmetry," J. Vision, Vol. 13, 18, 2013.

[10] M. E. Rudd, "A cortical edge-integration model of object-based lightness computation that explains effects of spatial context and individual differences," Front. Hum. Neurosci., Vol. 8, 640, 2014.

[11] A. Gilchrist, C. Kossyfidis, F. Bonato, T. Agostini, J. Cataliotti, X. Li, B. Spehar, V. Annan, & E. Economou, "An anchoring theory of lightness perception," Psychol. Rev., Vol. 106, 795-834, 1999.

[12] Chevreul, M. E., The principles of harmony and contrast of colors an their applications to the arts. New York: Van Nostrand Reinhold, 1839/1970.

[13] R. Shapley & R. C. Reid, "Contrast and assimilation in the perception of brightness," Proc. Natl. Acad. Sci. USA, Vol. 82, 5983-5986, 1985.

[14] R. C., Reid & R. Shapley, "Brightness induction by local contrast and the spatial dependence of assimilation," Vision Res., Vol. 28, 115-132, 1988.

[15] M. E. Rudd & I. K. Zemach, "Contrast polarity and edge integration in achromatic color perception," J. Opt. Soc. Am. A, Vol. 24, 2134-2156, 2007.

[16] W. von Bezold, The theory of color in its relation to art and art-industry. Boston: L. Prang & Co, 1876.

[17] H. Helson, "Studies of anomalous contrast and assimilation," J. Opt. Soc. Am., Vol. 53, 179-184, 1963.

[18] M. E. Rudd & D. Popa. "Stevens' brightness law, contrast gain

control, and edge integration in achromatic color perception: A unified model," J. Opt. Soc. Am. A, Vol. 24, 2766-2782, 2007.

[19] M. E. Rudd, "Lightness computation by a neural filling-in mechanism," Proc. SPIE, Vol. 4299, 400-413, 2001.

[20] M. E. Rudd, & K. F. Arrington, "Darkness filling-in: A neural model of darkness induction," Vision Res., Vol. 41, 3649-3662, 2001.

[21] D. H, Hubel & T. N, Wiesel, "Receptive fields of single neurons in the cat's striate cortex," J. Physiol. (Lond.), Vol. 148, 574-591, 1959.

[22] D. H. Hubel & T. N, Wiesel, "Receptive fields, binocular interaction, and functional architecture of monkey striate cortex," J. Physiol. (Lond.), Vol. 160, 106-154, 1962

[23] D. H. Hubel, Eye, Brain, and Vision. New York: W. H. Freeman, 1988.

[24] S. Marcelja, "Mathematical description of the responses of simple cortical cells," J. Opt. Soc. Am., Vol. 70, 1297-1300, 1980.

[25] J. G. Daugman, "Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters." J. Opt. Soc. Am. A, Vol. 1, 1160-1169, 1985.

[26] J. Kremkow, J. Jin, S. J. Komban, Y. Wang, R. Lashgari, X. Li, M. Jansen, Q. Zaidi, & J-M. Alonso, "Neuronal nonlinearity explains greater visual spatial resolution for darks than lights," Proc, Natl. Acad. Sci. USA, Vol. 111, 3170-3175, 2014.

[27] A. Galmonte, A. Soranzo, A., M. E. Rudd, & T. Agostini, "The phantom illusion," Vision Res., Vol. 117, 49-58, 2015.

[28] H. L. F. von Helmholtz, Treatise on Physiological Optics. J. P. C. Southall (Transl.). Rochester, NY: Opt. Soc. Am.

[29] B. N. Bushnell, P. J. Harding, Y. Kosai, W. Bair, & A. Pasupathy, "Equiluminance cells in visual cortical area V4," J. Neurosci. Vol. 31, 12398–12412, 2011.

[30] C. Kennard, M. Lawden, A, B, Morland, & K. H. Ruddock, "Color discrimination and color constancy are impaired in a patient with incomplete achromatopsia associated with prestriate cortical-lesions," Proc. Roy. Soc. Lond. B, Vol. 206, 169-175, 1995.

[31] S, Zeki & L. Marini, L, "Three cortical stages of colour processing in the human brain," Brain, Vol. 121, 1669-1686, 1998.

[32] R. W. Kentridge, C. A. Heywood, & A., "Chromatic edges, surfaces and constancies in cerebral achromatopsia," Neuropsychologia, Vol. 42, 821-830, 2004.

[33] C. D. Gilbert & W. Li, "Top-down influences on visual processing," Nat. Rev. Neurosci., Vol. 14, 350-363, 2013; doi: 10.1038/nrn3476

Author Biography

Michael Rudd earned BS degrees in physics and psychology from the University of California, Davis (1978) and a PhD in experimental psychology from the University of California, Irvine (1987). He has taught on the faculties at Johns Hopkins University and the University of Washington; and has worked as a Research Scientist for the Howard Hughes Medical Institute. His current research straddles computational neuroscience and psychophysical investigations of lightness, color, and human rod vision.