Psychological Review
1999, Vol. 106, No. 4, 795–834

0033-295X/99/$3.00

An Anchoring Theory of Lightness Perception

Alan Gilchrist and Christos Kossyfidis
Rutgers, The State University of New Jersey,
Newark College of Arts and Sciences

Frederick Bonato
St. Peters College

Tiziano Agostini
University of Trieste

Joseph Cataliotti
Ramapo College

Xiaojun Li
Lucent Technologies

Branka Spehar
University of New South Wales

Vidal Annan and Elias Economou
Rutgers, The State University of New Jersey, Newark College of Arts and Sciences

A review of the field of lightness perception from Helmholtz to the present shows the most adequate theories of lightness perception to be the intrinsic image models. Nevertheless, these models fail on 2 important counts: They contain no anchoring rule, and they fail to account for the pattern of errors in surface lightness. Recent work on both the anchoring problem and the problem of errors has produced a new model of lightness perception, one that is qualitatively different from the intrinsic image models. The new model, which is based on a combination of local and global anchoring of lightness values, appears to provide an unprecedented account of a wide range of empirical results, both classical and recent, especially the pattern of errors. It provides a unified account of both illumination-dependent failures of constancy and background-dependent failures of constancy, resolving a number of long-standing puzzles.

We present a new theory of how the visual system assigns lightness values—or perceived black, white, or gray values—to various regions of the retinal image. The theory grew out of two problems facing what are probably the most advanced models of lightness perception: the intrinsic image models.

Alan Gilchrist, Christos Kossyfidis, Vidal Annan, and Elias Economou, Psychology Department, Rutgers, The State University of New Jersey, Newark College of Arts and Sciences; Frederick Bonato, Department of Psychology, St. Peters College; Tiziano Agostini, Department of Psychology, University of Trieste, Trieste, Italy; Joseph Cataliotti, Department of Psychology, Ramapo College; Xiaojun Li, Lucent Technologies, Warren, New Jersey; Branka Spehar, Department of Psychology, University of New South Wales, Sydney, New South Wales, Australia.

Tragically, Christos Kossyfidis passed away in March 1996. Among other contributions to this article, he produced evidence that lightness is anchored solely at white, not at both white and black.

We acknowledge the support of National Science Foundation Grants BNS-8909182, DBS-9222104, and SBR 95-14679.

We thank the following people who commented on an earlier version of this article: Larry Arend, Nicola Bruno, Dave Emmerich, Alesandra Galmonte, Walter Gerbino, Holger Knau, Luiz Pessoa, William Ross, Hal Sedgwick, Dejan Todorović, and Paul Whittle.

Correspondence concerning this article should be addressed to Alan Gilchrist, Psychology Department, Rutgers, The State University of New Jersey, Newark College of Arts and Sciences, Newark, New Jersey 07102.

Theoretical Developments: From Unconscious Inference to Intrinsic Image

Basic Ambiguity: Luminance Versus Lightness; Lightness Constancy

To understand the problem of lightness constancy, it is necessary to understand the ambiguous relationship between luminance values (light intensities) within the retinal image and the lightness values of the surfaces in our perceived world. A key problem is that luminance values in the retinal image are a product, not only of the actual physical shade of gray of the imaged surfaces, but also, and even more so, of the intensity of the light illuminating those surfaces. The luminance of any region of the retinal image can vary by a factor of no more than thirty to one as a function of the physical reflectance of that surface. However, it can vary as a factor of a billion to one as a function of the amount of illumination on that surface. The net result is that any given luminance value can be perceived as literally any shade of gray, depending on its context within the image. Despite this challenge, people perceive shades of surface grays with rough accuracy. This is the well-known problem of lightness constancy. Here is a brief history of attempts to solve this problem.

Inferring the Illuminance: Helmholtz

In one of the earliest attempts to solve this problem, Helmholtz (1866) proposed that the luminance of a region in the retinal image is compared with the perceived intensity of the illumination in that region of the visual scene. Dividing the luminance by the illumination yields reflectance, and this is how a physicist might determine the reflectance of a surface. However, as an account of human lightness, the theory has been fraught with both logical and empirical difficulties.

Hering's Paradox

Hering (1874/1964) argued that Helmholtz's (1866) position is circular. Given the luminance of a surface, one would need to know its reflectance to infer the illumination, but reflectance is what one is trying to find. Hering emphasized the role of sensory mechanisms like pupil size, adaptation, and lateral inhibition. However, he also attributed constancy to a cognitive variable he called memory color.

The publication of Katz's (1911, 1935) book *The World of Colour* gave enormous momentum to the emerging field of lightness perception. "Its importance at the time of its publication can hardly be overrated," wrote Koffka (1935, p. 241). Katz presented a thorough phenomenological analysis of the visual experience of color. He outlined the various modes of color appearance, emphasizing especially the distinction between surface color and film color. Yet at the same time, Katz was a rigorous experimentalist. He developed a variety of experimental methods for studying lightness constancy, including the now standard method involving side-by-side fields of light and shadow. He showed that lightness constancy holds even when all three of Hering's (1874/1964) mechanisms plus memory color are ruled out.

Gestalt and the Relational Approach

The gestalt theorists rejected the assumption that light per se is the stimulus for lightness, in favor of luminance ratios or gradients. Koffka (1935) wrote, "Our theory of whiteness constancy will be based on this characteristic of colours . . . that perceived qualities depend upon stimulus gradients" (p. 244). Gelb (1929) showed that a piece of black paper appears white when presented alone in a spotlight but much darker when a real white is placed next to it in the spotlight. Experiments of fundamental importance were also conducted by Benary (1924), Kardos (1934, 1935), Burzlaff (1931), Wolff (1933, 1934), and Katona (1935), to name only a few. Reading the current literature, one would hardly suspect the vigorous empirical and theoretical developments that took place during the several decades following the publication of Katz's book in 1911.

Those developments were derailed by the events leading to World War II. After the war, the spotlight shifted to America, where contrast theories, spurred by the first direct physiological evidence for lateral inhibition, came to dominate the field completely. The stimulus conditions became highly reduced: luminous patches presented in dark rooms. Relative luminance came to mean contrast. The gestalt lessons were lost in the stampede to explain lightness at the physiological level. The contrast theorists argued that the facts of lateral inhibition rendered the vague ideas of the gestaltists obsolete. Consequently, they felt no need to cite the earlier European work, and they did not (see Gilchrist, 1996).

Several important publications in the tradition of relational determination appeared in the 1940s, but they were quickly assimilated to the contrast interpretation. Helson (1943, 1964), like Koffka, based lightness values on stimulus gradients. He proposed that the luminance of a target surface is compared with the average luminance (in fact, a weighted average) in the retinal image, such that a surface with a luminance equal to the average luminance is seen as middle gray, luminances above the average are seen as light gray or white, and those below the average are seen as dark gray or black.

Wallach (1948), in a landmark article, proposed a simple ratio theory of lightness. Presenting observers with two disk–annulus displays, he showed that disks of different luminance appear equal in lightness as long as the disk–annulus luminance ratios are equal.

Intrinsic Image Theories

Empirical evidence accumulated more recently (Shapley & Enroth-Cugell, 1984; Whittle & Challands, 1969; Yarbus, 1967) tends to support the idea that retinal encoding processes simply encode relative luminance (see Gilchrist, 1994). Yet, there has been a reluctance to embrace such a simple idea.

In the past 2 decades, more effective models of lightness perception have emerged. They have evolved, not out of contrast theories, but out of attempts to correct several limitations in Wallach's (1948) simple ratio formula. Although it has become widely agreed that the concept of luminance ratios at edges goes far toward explaining the traditional problem of lightness constancy, this same insight has revealed a second constancy problem. When the same piece of gray paper is viewed successively against different backgrounds, the luminance ratio at the edge of the paper changes dramatically, yet the paper appears to change very little in lightness, contrary to the ratio principle. This constancy of lightness with respect to changing background has been labeled *Type II constancy,* to distinguish it from lightness constancy with respect to changing illumination, or *Type I constancy.* W. Ross and L. Pessoa (personal communication, December 20, 1996) have proposed the more memorable terms *illumination-independent constancy* (Type I) and *background-independent constancy* (Type II), and we have adopted that usage. A third kind might be called veil-independent constancy (Gilchrist & Jacobsen, 1983).

Edge integration. Wallach's (1948) ratio rule deals effectively with adjacent luminances. However, background-independent constancy seems to require a mechanism by which the luminance values of widely separated regions in the retinal image can be compared. The recognition of this need led to several articles, which appeared at almost the same time. Land and McCann (1971); Arend, Buehler, and Lockhead (1971); and Whittle and Challands (1969) offered evidence and arguments suggesting that the visual system is capable of deriving the luminance ratio between two surfaces remote from each other in the image. The exact mechanism for this is unknown, but one suggestion is that luminance ratios at every edge encountered along an arbitrary path from one surface to its remote pair are mathematically integrated.

Edge classification. Gilchrist (1977, 1979, 1980) demonstrated empirically an observation that had been made by Koffka in 1935 that "not all gradients are equally effective as regards the appearance of a particular field part" (p. 248) and, in another passage, that "clearly two parts at the same apparent distance will, ceteris paribus, belong more closely together than field parts organized in different planes" (p. 246). Gilchrist found that perceived lightness values depend primarily on luminance ratios between adjacent regions perceived to lie in the same plane, as opposed to luminance ratios between any two adjacent parts of the visual field.

Inspired by Koffka's (1935) observation that some luminance ratios are relatively effective in determining surface lightness whereas others are relatively ineffective, Gilchrist (1977, 1979; Gilchrist, Delman, & Jacobsen, 1983) proposed a distinction between what he called reflectance edges and illuminance edges. Reflectance edges are those luminance borders in the retinal image that are caused by a change in the reflectance (or pigment) of the surface being viewed, whereas illuminance edges are those that are caused by changes in the intensity of the illumination on a surface, such as the border of a cast shadow, for example, or the luminance step at a corner. Gilchrist proposed that the visual system must classify edges in the image into one of these two main categories, before edge integration. Then, an integration of all the edges in the reflectance category can yield a map of all the reflectances in the visual field, just as an integration of all the edges in the illuminance class yields a map of the illuminance pattern within the visual field. In effect, Gilchrist proposed to use edge classification as a wedge to pry the retinal image into two overlapping layers, one representing surface lightness values, the other representing the pattern of illuminance on those surfaces.

Parsing into layers. At the same time, Bergström (1977) proposed that luminance variations within the retinal image are vector analyzed into three components: one for surface reflectance, one for illumination, and one for three-dimensional form. Bergström's distinction between illumination changes and three-dimensional form changes is roughly equivalent to Gilchrist's further breakdown of illumination edges into cast illuminance edges and attached illuminance edges.

Adelson and Pentland (1996) have recently proposed an elegant scheme that bears a striking resemblance to Bergström's (1977) model. They likened the visual system to a three-person workshop crew that produces theatrical sets. One person is a painter, one is an illumination expert, and one bends metal. Any luminance pattern can be produced by any of the three specialists. The painter can paint the pattern. The lighting expert can produce the pattern with variations in the illumination, and the metalworker can create the pattern by shaping the surface, as in shape from shading. But, for a given luminance pattern, the cost of these three methods is not the same, setting up an economy principle in which each desired pattern is to be created in the cheapest possible way, reminiscent of the gestalt simplicity principle (Gerbino, 1994).

In 1978, Barrow and Tenenbaum proposed that every retinal image is composed of a set of what they called intrinsic images. One intrinsic image would contain the array of reflectance values in the scene; a second, the array of illumination intensities; a third, the array of depth values; and so on.

The Bergström (1977), Gilchrist (1979), Barrow and Tenenbaum (1978), and Adelson and Pentland (1996) models all have in common the idea that the retinal image is parsed into a set of overlapping layers, much as in the scission idea made popular by Metelli (1985; Metelli, Da Pos, & Cavedon, 1985) in his analysis of perceived transparency. We refer to these models as intrinsic image models. An excellent discussion of them can be found in a chapter by Arend (1994).

Two Weaknesses of the Intrinsic Image Models

Intrinsic image models are the most advanced models in the continuing development of lightness theory. They offer an explanation of both illumination-independent and background-independent constancy. Yet, they are incomplete in two very important, though different, ways: (a) They cannot account for the pattern of errors in lightness perception, and (b) they have no anchoring rule. We consider these in turn.

The problem of errors. The goal of the computational enterprise that produced the intrinsic image models has been the modeling of a completely veridical lightness perception system. This is implied in terms like *inverse optics* and *recovering reflectance.* In that sense, it has been consistent with the goals of machine vision. Failures of constancy and other perceptual errors have been largely ignored. Even in human vision there is a good reason for this approach. The achievement of constancy and veridicality in the perception of surface lightness is stunning, especially when the various challenges to constancy are recognized. This degree of veridicality does not happen by accident; it cannot be merely the by-product of a system with goals other than veridicality. Somehow a very robust truth-seeking quality must lie at the heart of the system. Thus, if the achievement of veridicality is considered to be more impressive than the degree of failure, it makes sense to begin by trying to model a system that comes as close to veridical perception as possible.

Nevertheless, unlike the situation in machine vision, a theory of human lightness must include an account of errors. To the extent that veridicality fails in lightness perception and to the extent that the intrinsic image models predict veridicality, the models must fail to account for human lightness perception. But perhaps more importantly, a systematic analysis of human lightness errors can be a powerful tool for revealing how surface lightness is processed by the human visual system. There is a simple and compelling logic behind the study of errors.

1. Errors in lightness perception are always present, however slight.
2. These errors are systematic, not random.
3. The pattern of errors must reflect visual processing; it must be the signature of the visual system.

The theoretical picture would be brighter if the empirical pattern of errors could be produced by tweaking the veridicality models. However, no such prospect is in sight. For example, there is no coherent approach that can explain both illumination-dependent failures of constancy (Type I) and background-dependent failures of constancy (Type II), despite an attempt by Gilchrist (1988). Background-dependent failures of constancy, of which the textbook version of simultaneous lightness contrast is the best known example, are typically described as simply reflecting the operation of the mechanism that achieves illumination-independent constancy. But how are illumination-dependent failures of constancy to be explained? Is there no systematic relationship between these two kinds of failure? We consider a unified account of illumination-dependent and background-dependent failures to be a major goal of a theory of lightness errors.

The missing anchor. The second shortcoming in the intrinsic image models, in their current form, is that they are missing an essential component for veridical perception: an anchoring rule.

We now turn to an extended discussion of the anchoring problem, which in turn brings us back to the errors problem.

We find that the attempt to fill a crucial gap in the intrinsic image approach by finding the missing anchoring rule in fact turns out to undermine the entire intrinsic image approach. At the same time, it provides the outlines of a very different approach to surface lightness, one that excels in its ability to explain the pattern of errors found across a very broad range of empirical results.

The Anchoring Problem

A Concrete Example

Although the ambiguous relationship between the luminance of a surface and its perceived lightness is widely understood, there has been little appreciation of the fact that relative luminance values are scarcely less ambiguous than absolute luminance values. For instance, consider a pair of adjacent regions in the retinal image whose luminance values stand in a 5:1 ratio. This 5:1 ratio informs the visual system only about the relative lightness values of the two surfaces, not their specific or absolute lightness values. It informs only about the distance between the two gray shades on the phenomenal gray scale, not the specific location of either on that scale. There is an infinite family of pairs of gray shades that are consistent with the 5:1 ratio. For example, if the 5 represents white, then the 1 represents middle gray. However, the 5 might represent middle gray, in which case the 1 would represent black. Indeed, it is even possible that the 1 represents white and the 5 represents an adjacent self-luminous region. So, the solution is not even restricted to the scale of surface grays.

The anchoring problem is the problem of how the visual system ties relative luminance values extracted from the retinal image to specific values of perceived black, white, and gray.

Mapping Luminance Onto Lightness

To derive specific shades of gray from relative luminance values in the image, one needs an anchoring rule. An anchoring rule defines at least one point of contact between luminance values in the image and gray scale values along our phenomenal black-to-white scale.[1] Lightness values cannot be tied to absolute luminance values because there is no systematic relationship between absolute luminance and surface reflectance, as noted earlier. Rather, lightness values must be tied to some measure of relative luminance.

The anchoring problem, closely related to the issue of normalization (Horn, 1977, 1986), is illustrated in Figure 1. Consider the following remark by Koffka (1935) in relation to Figure 1:

> The stimulus gradient . . . alone does not determine the absolute position of this apparent gradient . . . . This whole manifold of colours may be considered as a fixed scale, on which the two colours produced by the two stimulations . . . keeping the same distance from each other, may slide, according to the general conditions. I have called this the principle of the shift of level. (p. 255)

Anchoring concerns where this sliding relationship comes to rest. Koffka suggested that the relative lightness of two regions in an image can remain fully consistent with the luminance ratio between them, even though their absolute lightness levels depend on how the luminance values are anchored.

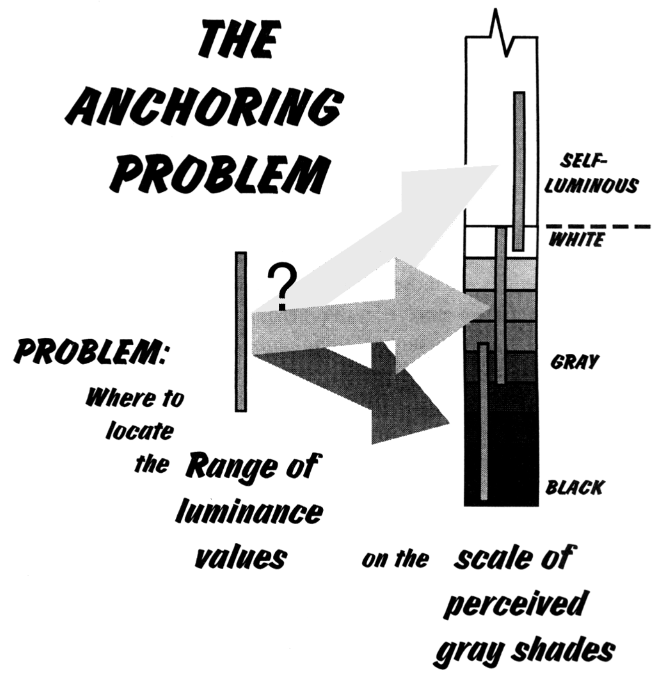

Figure 1. The anchoring problem. From "Relative Area and Relative Luminance Combine to Anchor Surface Lightness Values," by X. Li and A. Gilchrist, 1999, *Perception & Psychophysics, 61,* p. 772. Copyright 1999 by X. Li and A. Gilchrist. Reprinted with permission.

Wallach on Anchoring: Highest Luminance Rule

Apparently, Wallach made no systematic study of the anchoring problem. He did mention in passing that the value of white is assigned to the highest luminance in the display and serves as the standard for darker surfaces (Wallach, 1976, p. 8). Land, McCann, and Horn (Horn, 1977; Land & McCann, 1971; McCann, 1987, 1994) also adopted this rule, which we call the highest luminance rule.

Helson on Anchoring: Average Luminance Rule

Although Helson (1943) neither discussed the anchoring problem per se nor made any systematic investigation of it, his entire adaptation-level theory is built, in effect, on an anchoring rule. His rule is that the average luminance in the visual field is perceived as middle gray, and this value serves as the standard for both lighter and darker values. This *average luminance rule,* which is closely related to the gray world hypothesis (Hurlbert, 1986), has found its way into more recent models (Buchsbaum, 1980), especially in the chromatic domain. It is implicit in the concept of equivalent background (Brown, 1994; Bruno, 1992; Bruno, Bernardis, & Schirillo, 1997; Schirillo & Shevell, 1996). Land (1983) reverted to the average luminance rule in a later version of Retinex theory.

[1] If the relative luminances in the image are anchored at a single point, then the lightness value of each surface can be determined by the luminance ratio between that surface and the anchor. We refer to this part of the process as scaling.

No one has in fact proposed a rule whereby the lowest luminance in the field is seen as black, but in principle such a rule is possible.

Anchoring in Intrinsic Image Models

None of the intrinsic image models contain any specific anchoring rule. However, the task is simplified because the various factors that can modulate luminance within the image—that is, reflectance, illuminance, and three-dimensional form—have already been segregated into separate layers. The obvious next step would be to apply something like the highest luminance rule solely to the reflectance intrinsic image.

Anchoring Versus Scaling

Anchoring concerns the mapping of relative luminance values from the image onto perceived shades of surface gray. A parallel aspect of this mapping, which we call the *scaling problem,* is illustrated in Figure 2. It is distinct from the anchoring problem, though possibly less important. It concerns how the range of luminances in the image is mapped onto a range of perceived grays. Rescaling can take one of two forms: compression, in which the range of perceived surface grays is less (on a log scale) than the range of luminances in the image, and expansion (Brown & MacLeod, 1992, used the term *gamut expansion*), which is just the reverse.[2]

Although anchoring can be said to concern the constancy of absolute shades of gray, scaling can be said to concern the constancy of relative shades of gray. Wallach (1948) was more explicit about scaling than about anchoring. His ratio principle is a scaling rule; it asserts a one-to-one mapping from relative luminance to relative lightness.

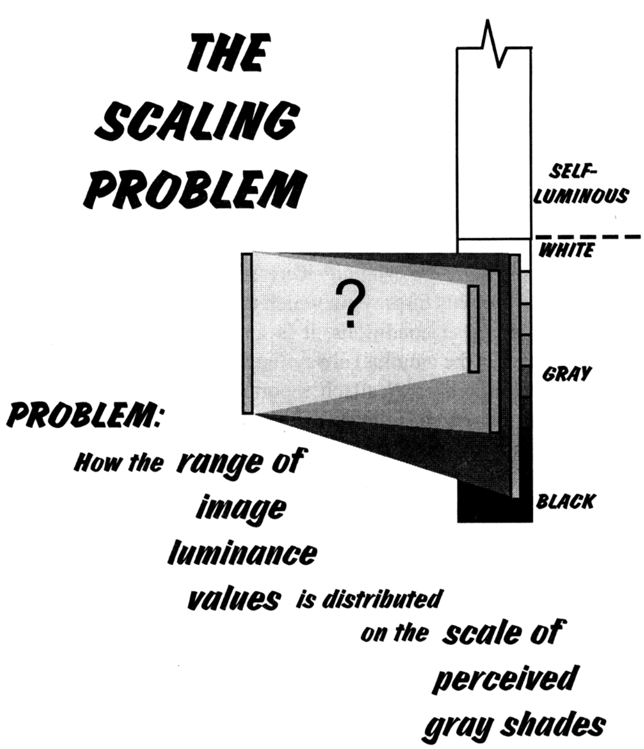

Figure 2. The scaling problem concerns constancy for the range of perceived grays. The perceived range of grays can be expanded, compressed, or equal to the range of luminances in the stimulus.

Contrast theories, by comparison, reject this kind of one-to-one mapping in favor of expansion. In contrast theories, the relative luminance values encoded at input are magnified (Hurvich & Jameson, 1966, p. 85), exaggerated (Leibowitz, 1965, p. 57), or amplified (Cornsweet, 1970, p. 300) by the lateral inhibitory component of the encoding process itself.

Helson (1943), unlike Wallach (1948), was more explicit about anchoring than about scaling. His adaptation level represents an anchor; he did not specify how regions lighter and darker than the average are scaled, but we presume that the ratio principle would apply.

The Rules of Anchoring in Simple Displays

Our approach is to consider the rules of anchoring under minimum conditions for the perception of a surface and then to attempt to describe how the rules change as one moves systematically from simple images to complex images. We find that, for simple images, anchoring depends on two dimensions of the stimulus: relative luminance and relative area.

What Are Minimal Conditions?

Katz (1935) and Gelb (1929) were among the first to observe that the visual perception of a surface requires the presence of at least two adjacent regions of nonzero luminance. Wallach (1963) noted, "Opaque colors which deserve to be called white or gray, in other words 'surface colors,' will make their appearance only when two regions of different light intensity are in contact with each other" (p. 112). Ideally, these minimal conditions would be met by a pair of surfaces that fill the entire visual field (Koffka, 1935, p. 111). One can, for instance, use the inside of a large hemisphere painted in two gray shades to cover an observer's entire visual field. Heinemann (1972) expressed a common view that the "simplest experimental arrangement for studying induction effects" (p. 146) consists of two surfaces presented within a void of total darkness. In our view, however, these conditions only approximate the simplest because there are at least three regions in the visual field, including the dark surround, instead of two. There are two borders rather than the minimum of one.

Highest Luminance Versus Average Luminance

Which rule is correct under such minimal conditions: Wallach's (1948) highest-luminance-as-white rule or Helson's (1943) average-luminance-as-middle-gray rule? Predictions made by the two rules diverge when the stimulus array contains a range of luminances substantially less than the 30:1 range between white and black.

We placed observers' heads inside a large acrylic hemisphere, the inside surface of which was divided into two halves of equal

[2] We typically use the term *perceived surface gray* or *perceived reflectance* to mean the reflectance of the chip from the Munsell scale that the observer perceives to be the same lightness as the target surface.

size (Li & Gilchrist, 1999). One half was painted completely matte black (luminance: 1.92 cd/m^2); the other half was painted middle gray, as shown in Figure 3a. The middle-gray half was seen as fully white, and the black half was seen as a darkish middle gray (luminance: 10.6 cd/m^2). All observers reported this result, and the variability was very low. This outcome decisively favors the highest luminance rule over the average luminance rule.

Other findings, using more complex stimuli, agree with this conclusion. In a separate line of work, Gilchrist and Cataliotti (1994) presented observers with a flat Mondrian containing 15 rectilinear pieces of gray paper, spanning only a 10:1 reflectance range (from black to middle gray). The Mondrian was illuminated by a special projector and presented within a relatively darkened laboratory, and observers indicated their perceptions by making matches from a separately housed and lighted 16-step Munsell scale. In a further experiment (Cataliotti & Gilchrist, 1995), we placed observers' heads inside a small trapezoidal-shaped empty room, all the walls of which were covered with a Mondrian pattern spanning an even smaller range (4:1). In both of these Mondrian experiments, the highest luminance was perceived as white, and no black surfaces were seen.

The results for all three tests are shown in Figure 4. Notice that in all cases, the highest luminance is perceived as white, even if it is physically a dark gray surface. Notice further that the symmetry

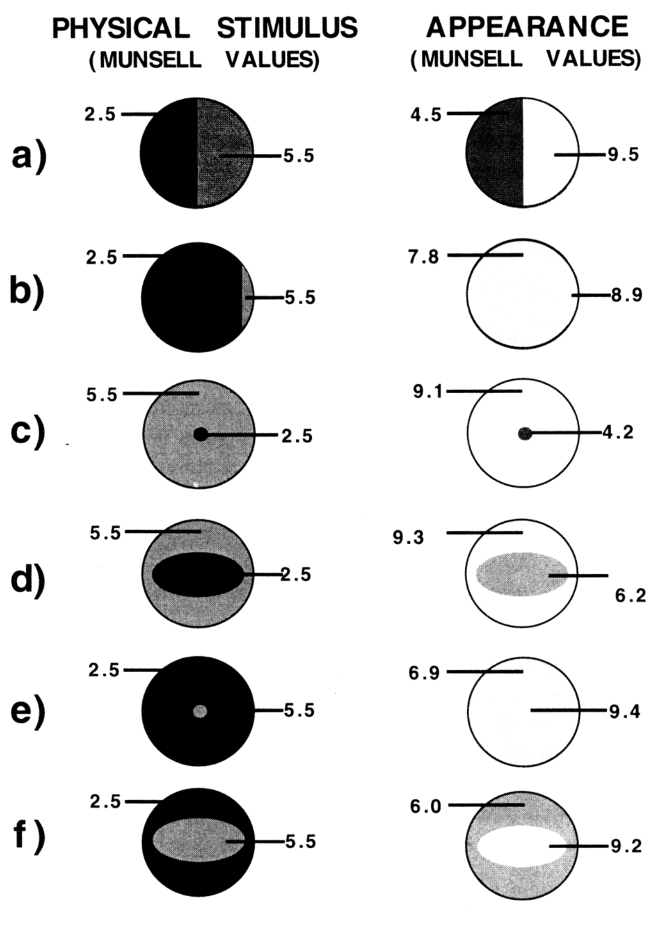

Figure 3. Actual and perceived gray shades in domes experiments (Li & Gilchrist, 1999).

implicit in the average luminance rule is missing from these results. Although a perceived white is always present in the scene, it is not always the case that a perceived black is present in the scene, and the perceived values do not distribute themselves symmetrically about middle gray.

The two rules have been tested indirectly in several experiments on equivalent surrounds. Bruno (1992) asked observers to make a brightness match between two target squares of equal luminance, one surrounded by a homogeneous region and one embedded in a checkerboard or Mondrian-like region. The luminance of the homogeneous surround that was chosen by the observers as having the same effect on a target as the checkerboard was a luminance value close to that of the highest luminance in the checkerboard, not the average.[3] Schirillo and Shevell (1996) presented a similar pair of displays to observers but asked them to make a brightness match between a target square on a checkerboard and one on a uniform background of the same average luminance, as the checkerboard contrast was increased from 0% to 100%. Their results for increments (regions with a luminance higher than that of the surround) agreed with those of Bruno, but their results for decrements showed a different pattern (see also Bruno et al., 1997).

McCann (1994) showed with chromatic Mondrians that, if the average color is held constant while the maximum excitation in a cone channel is varied, perceived colors in the Mondrian change. However, if the maximum excitation in a cone channel is held constant while the average is varied, perceived colors change only slightly (see also Brown, 1994).

Luminosity Problem: Direct Contradiction to the Highest Luminance Rule

There is one phenomenon, however, that directly contradicts the highest luminance rule: the perception of self-luminous surfaces (Bonato & Gilchrist, 1994; Ullman, 1976). According to the highest luminance rule, white is a ceiling; nothing can appear to be brighter than white.

Wallach on increments: The highest luminance rule does not apply. To be fair to Wallach (1948, 1976), his support for the highest luminance rule was neither emphatic nor unambiguous. Wallach's systematic experiments were conducted only with decrements: disk–annulus displays in which the disk is darker than the annulus. Under these conditions, it is an empirical fact that the higher luminance (the annulus) always appears white, regardless of its luminance. However, Wallach reported informally that when the disk is an increment relative to the annulus, it appears luminous. The implication is that the highest luminance rule applies to decrements but not to increments.

Common factor for decrements and increments: Geometric, not photometric. If the highest luminance rule can be applied only to decrements, what rule could apply to both decrements and increments? The common denominator appears to be a geometric factor, not a photometric factor: It is always the annulus that

[3] We believe that there is a modest expansion in the perceived relationship between the homogeneous background and its target that is associated with the limited range in that framework. This scale normalization effect, discussed in the section *Theoretical Value of Belongingness,* may explain why the matching value for the homogeneous background lies just below the maximum on the checkerboard.

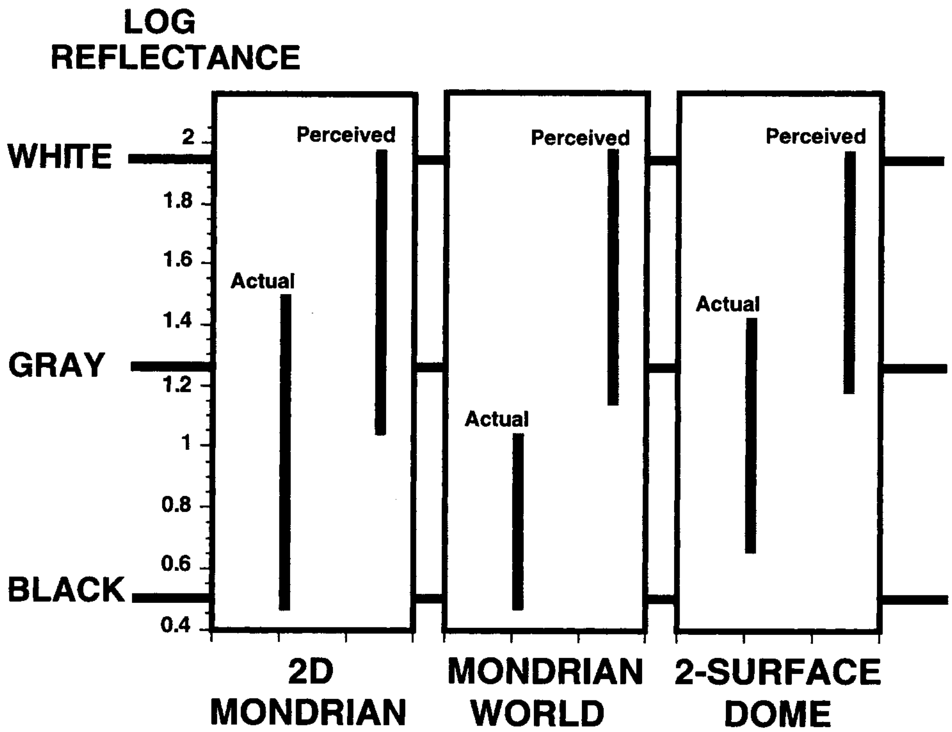

Figure 4. Evidence for the highest luminance rule, showing the actual range of grays and the perceived range in three separate experiments. 2D = two-dimensional.

appears white, both in decrements and in increments. So far, anchoring rules have been considered only in terms of luminance relations. This observation suggested that spatial relations play an important role.

Wallach's (1948) disk–annulus stimuli lend themselves readily to a figure–ground analysis, and Gilchrist and Bonato (1995) formulated the hypothesis that lightness values are anchored, not by any special luminance value, but by the surrounding or background region. The rule, which they called the surround rule, says that for simple displays, lightness is anchored by the surround, which always appears white. Gilchrist and Bonato tested the surround rule against the highest luminance rule using both a disk–annulus configuration and a disk–ganzfeld configuration. The disk–annulus experiment produced data roughly consistent with the surround rule but with some influence of the highest luminance rule. The disk–ganzfeld data were completely consistent with the surround rule.

Area Effects

In the disk–ganzfeld experiment, the figure–ground distinction was confounded with relative area. Using the dome method described earlier, Li and Gilchrist (1999) varied relative area while figure–ground arrangements were held constant. Schematic representations of these stimuli are shown in Figure 3, along with the results, given in Munsell values.

Surround Rule Fails

The net result is that the surround rule was completely undermined. In the incremental large oval condition, illustrated in Figure 3f, the background appeared middle gray (Munsell 6.0/), not white, as it should according to the surround rule. What had appeared to be a matter of figure–ground turns out to be a matter of relative area. This result is seen clearly in the split-dome conditions, shown in Figures 3a and 3b. As the darker region increases its area from half the visual field (Figure 3a) to most of the visual field (Figure 3b), its lightness moves from Munsell 4.5/ to 7.8/. This finding brings to mind Helson's (1964) comment that "we need assume only that within certain limits area acts like luminance, that is, increase in area has the same effect as increase in luminance" (p. 292).

Highest Luminance Plus Area

The highest luminance rule was rescued by combining it with a tendency for the largest area to appear white. Thus, anchoring under minimal conditions (two regions in the visual field) appears to depend on both photometric and geometric factors. We can say that the larger a surface, the lighter it appears, or "the larger the lighter" for short. However, the larger the lighter does not apply under all conditions. It applies only when there is a conflict between the tendency for the highest luminance to appear white and the tendency for the largest area to appear white. As long as the highest luminance has the largest area, there is no conflict, and that region becomes a very stable anchor. Darker surfaces are seen relative to that anchor simply according to the ratio principle, and the larger the lighter does not apply.

The Area Rule

By comparing our results with others in the literature, we were able to formalize a rule governing the effects of relative area on perceived lightness that had not been previously recognized. The rule, which we call the *area rule,* describes how relative area and relative luminance combine to anchor lightness perception. The rule is this: In a simple display, when the darker of the two regions has the greater relative area, as the darker region grows in area, its lightness value goes up in direct proportion.[4] At the same time the lighter region first appears white, then a fluorent white, and finally self-luminous.

Strictly speaking, the rule applies to visual fields composed of only two regions of nonzero luminance. Application of the rule to more complex images remains to be studied.

There are at least 10 other articles in the literature concerning the influence of area on lightness or brightness. We see that a great deal of order is brought to this part of the literature when these results are analyzed in terms of anchoring in general, and the area rule in particular.

Wallach (1948) tested three annulus-to-disk area ratios in his experiments with decremental disks: 4:1, 1:1, and 1:4. The first two of these yielded identical results. Relative area had an effect on perceived lightness only in the third case. When the disk, which was the darker of the two, had an area four times greater than that of the annulus, the disk appeared lighter than it would otherwise appear. The area rule predicts no difference between the 4:1 and 1:1 area ratios because the darker region (the disk) does not have the largest area. However, the rule does predict a difference between the 1:1 and 1:4 area ratios because the darker region does have the larger area. The darker region is predicted to lighten, and this is what Wallach reported.

Using lighting conditions similar to those used by Gelb (1929), Newson (1958) spotlighted a display consisting of a square target surrounded by a brighter square annular region. Holding both center and surround luminances constant, Newson tested perception of the center square while he varied the area of the surround from zero to an area roughly equal to that of the center square. This range is just the range within which the area rule applies. He obtained a pronounced effect on the lightness of the square. Moreover, his curve (Newson, 1958, Figure 4, p. 94) reaches an asymptote just where the areas of the center and surround become equal, suggesting that additional increases in the area of the surround would have no further effect on the lightness of the center.

Kozaki (1963) tested brightness using a haploscopic technique and a square center-surround display embedded in darkness. Because the area of the surround was always greater than the area of the test field, the area rule would apply under the conditions in which her test field was an increment. With increments, she obtained an area effect like that we obtained, consistent with our area rule.[5] However, she also obtained a weak area effect when the test field was a decrement.

Helson and his associates (Helson, 1963, 1964; Helson & Joy, 1962; Helson & Rohles, 1959) varied the relative area of either white or black stripes on a gray rectangle in an attempt to resolve the paradox of lightness contrast versus lightness assimilation, as in the classic von Bezold (1874) spreading effect. Their results are not consistent with the area rule. We are unable to resolve this discrepancy beyond the observation that the von Bezold effect may involve a relatively low-level kind of space-averaged luminance.

Burgh and Grindley (1962) reported no effect of area on perceived lightness using the traditional simultaneous lightness contrast display. However, it is crucial to note that they achieved their area changes by magnifying or minifying the entire display. As a consequence, the relative area between each gray target and its background was never changed, so this outcome does not contradict the area rule.

Yund and Armington (1975) also tested the dependence of brightness on relative area in a disk–annulus display. But, contrary to all the other studies, they tested the effect of the darker region on the brighter, an effect known to be either tiny or nonexistent (Freeman, 1967, p. 173; Heinemann, 1972).

We do not comment on two studies (Diamond, 1962; Whipple, Wallach, & Marshall, 1988) in which effects of area were studied because area was confounded with separation between test and inducing fields in those studies.

Four additional studies of brightness and area by Heinemann (1955), Diamond (1955), Stevens (1967), and Stewart (1959) are consistent with the area rule. These studies were part of the brightness induction literature and are considered in the section entitled Brightness Induction: Contrast or Anchoring? as part of a general review of that literature.

Luminosity and the Area Rule

Area rule implies luminosity. The surround rule had been proposed as a resolution of the apparent contradiction between the highest luminance rule and the perception of luminosity. If the surround rule must be abandoned, can the area rule resolve this contradiction and explain both surface lightness perception and luminosity? After all, Li and Gilchrist (1999) did obtain luminosity perception in several of their domes experiments even though their stimuli were nothing more than opaque surfaces. The area rule does appear to greatly illuminate the relationship between opaque surface lightness and self-luminosity. In general, we can say that luminosity perception occurs at one extreme end of the zone to which the area rule applies, that is, when a given surface is high in relative luminance but at the same time low in relative area.

Upward induction versus downward induction. When the luminance difference between two regions of the visual field is increased, two basic outcomes (or some combination) are possible. The darker region might remain perceptually constant while the lighter region moves toward and into a self-luminous appearance. This process might be called upward induction; it occurred in Gilchrist and Bonato's (1995) disk–ganzfeld experiments when the disk was an increment. Alternatively, the lighter region might remain perceptually constant while the darker region becomes

[4] The term *in direct proportion* here has the following meaning. When the darker region occupies 50% or less of the total area, its perceived lightness is determined simply by its luminance ratio with the lighter region according to Wallach's ratio principle and the highest luminance rule (the lighter region taken to be white). As the area of the darker region grows from 50% to 100% of total area, its lightness grows proportionately from this value toward white. We have empirical data supporting this claim regarding the lightness value at 50% and 100%, but we are making an assumption that the transition in lightness is smooth between these values.

[5] Although Figure 5 in Kozaki (1963) appears to show an area effect opposite to what we report, A. Kozaki has confirmed in a personal communication (September 5, 1995) that this is merely due to the omission of minus signs on the values along the abscissa that show "log ratio of test field area to inducing field area."

darker and darker gray. Such downward induction occurs in disk–annulus experiments (Heinemann, 1955; Wallach, 1948) when the disk is a decrement.

The question then becomes, when the luminance difference between two regions is increased, what determines whether this increased difference is experienced as an induction of luminosity into the brighter region or grayness into the darker region? This, of course, is another way of stating the anchoring problem. Indeed Schouten and Blommaert (1995a) have put the problem in this way.

In short, the answer appears to lie in relative area. When the darker region is large relative to the lighter region, most of the effect is expressed as upward induction in the lighter region, with only a small amount of downward induction. However, when the area of the darker region is small relative to the lighter region, there is very little upward induction; most of the effect is expressed as downward induction.

Figure 5 represents what is currently our best understanding of the appearance of the two regions in a simple framework as the relative area shifts from the lighter region to the darker. Typically, a figure of this kind would plot perceived lightness and luminosity as a function of luminance ratio, with relative area held constant. Note that in this figure we have plotted lightness and perceived luminosity as a function of relative area, with luminance ratio held constant! We can understand this graph by walking through it, moving from left to right as the x-axis shows increasing relative area of the darker region. Beginning with the dark–light border at the extreme left eccentric position, the darker region is very small relative to the lighter region. In this case, the lighter region will appear white, and the lightness of the darker region will depend simply on its luminance ratio with the lighter region. As the border shifts from the extreme left eccentric position to the center position, no change will occur in the perception of either region because all of these stimuli lie outside the zone of applicability of the area rule. Only when the border crosses the midpoint and begins to move toward the right-hand eccentric position does the area rule apply and begin to produce perceptual changes.

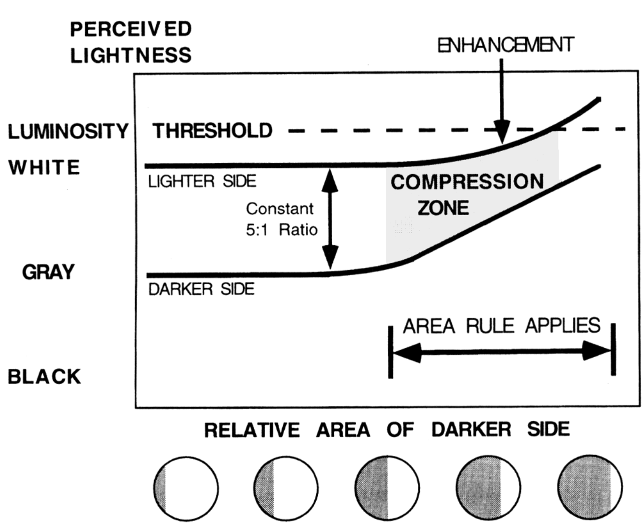

Figure 5. Schematic showing how lightness varies as a function of relative area in a two-field dome with constant luminance values. From "Relative Area and Relative Luminance Combine to Anchor Surface Lightness Values," by X. Li and A. Gilchrist, 1999, *Perception & Psychophysics, 61,* p. 782. Copyright 1999 by X. Li and A. Gilchrist. Reprinted with permission.

As the border passes the midpoint, the darker region begins to grow lighter and lighter. The lighter region continues to appear white, gradually becoming a more fluorent white (Evans, 1959, 1964, 1974), despite the constant luminance ratio. However, the lightening of the darker region is a stronger effect than the brightening of the lighter region. For example, with a luminance ratio of 30:1, the lightness of the darker region can move all the way from black to white (the lightness of the darker region approaches white as its relative area approaches 100%), while the lighter region remains at white. This implies a perceptual compression.[6] That is, the difference between the perceived lightness of the darker region and the perceived lightness of the lighter region is reduced even though the physical difference remains constant. The perceived lightness values seem literally to be squeezed between the tendency of the lighter region to appear white by the highest luminance rule and the tendency of the darker region to appear white because of its preponderant area.

Phenomena related to the area rule. Several additional effects are associated with this compression, or squeezing. One seems to be an enhancement of the lighter region, causing it to appear as a kind of preluminous super white (Heinemann, 1955; MacLeod, 1947). This twilight zone between white and luminosity has been termed *fluorence* by Evans (1959, 1974). We believe this may be the same enhancement phenomenon that Heinemann reported to occur in his test disk when the luminance of the annulus was increased while its luminance was still below that of the test disk. As the area of the darker side approaches 100%, its perceived lightness approaches white, and as this happens, the lighter region is forced to relinquish its white appearance (with its opaque quality) and take on the appearance of self-luminosity.

Schouten and Blommaert (1995a, 1995b) recently reported a phenomenon that they described as a novel compression mechanism in the luminance-brightness mapping. They called it brightness indention. Using a display that consisted of two disks within a ganzfeld, Schouten and Blommaert found that when both disks are brighter than the ganzfeld, the ganzfeld background does not appear homogeneous; it appears darker in the immediate vicinity of the disks, creating a kind of dark halo around each. Newson (1958, p. 87) described what appears to be the same phenomenon. This phenomenon occurs only in the zone to which our area rule applies, when the ganzfeld, with its large area, is darker than both of the disks. We believe this happens for the same reason as fluorence and the enhancement effect, namely because of the competition between the tendency for the ganzfeld to appear white because it has the greatest area and the tendency for the two disks to appear white because they have the highest luminance. Apparently, in this case the conflict is resolved by sacrificing the perceived homogeneity of the ganzfeld.

[6] This analysis might be considered unfair as the lighter region is prevented from becoming lighter because of a ceiling effect. However, the same compression happens when displays of three or four regions are tested. Compression will occur for the two or three darkest regions, even though none of these are subject to the ceiling effect.

Luminosity threshold and area: Bonato and Gilchrist. Bonato and Gilchrist (1994) studied luminosity thresholds by measuring the luminance value at which a target surface begins to appear self-luminous. These experiments were subsequently replicated (Bonato & Gilchrist, 1999) with larger targets. This produced higher thresholds. A 17-fold increase in the area of the target produced a 3-fold increase in the luminance required for luminosity perception. A corollary result is that as the luminance of the large target was increased, the increasing luminance ratio between the target and its background showed up as a darkening of the perceived surface lightness of the background. These results are consistent with the area rule.

Maximum luminance is not the same as the anchor. In view of the role of relative area in anchoring, it seems no longer appropriate to use the terms *anchor* and *highest luminance* interchangeably. The anchor in a given framework is the luminance value that appears white, and because of the area effect, the luminance value that appears white is not necessarily the highest luminance. The highest luminance is the same thing as the anchor only when the area rule does not apply.

Anchoring Rules for Simple Images: A Summary

For simple images, the rules of anchoring can be stated very concretely. Except when the area rule is engaged, anchoring is straightforward: The brightest region appears white, and the appearance of each darker region depends on its relationship to the white region, according to the formula

$$PR = L_t/L_h \times 90\%, \tag{1}$$

where PR is perceived reflectance, L_t is the luminance of the target, L_h is the highest luminance in the framework, and 90% is the reflectance of white.[7]

For conditions to which the area rule applies, the formula must be modified. Although greater precision will have to come from additional research, the evidence we have as of now, which has been expressed in graphic form in Figure 5, can be summarized algebraically as follows:

$$PR = (100 - A_d)/50 \times (L_t/L_h \times 90\%) + (A_d - 50)/50 \times (90\%), \tag{2}$$

where A_d is the area of the darker region, as a percentage of the total area in the field. The formula simply says that if A_d is 50% of the total area, the perceived reflectance of the darker region is just as it is given in Equation 1. As A_d approaches 100%, its perceived reflectance approaches 90%. Between these two endpoints, there is a smooth transition. The lighter region has no lightness value other than white, but as A_d grows, the lighter region takes on additional qualities, first fluorence, and finally self-luminosity. This qualitative change, not surprisingly, is difficult to capture mathematically.

Anchoring in Complex Images: A New Theory

So far the anchoring approach has been applied only to simple retinal images, and here it has proven its effectiveness in a compelling way. However, the ultimate challenge for a lightness theory lies in the kind of complex images we encounter routinely. How can the rules of anchoring by relative luminance and by relative area be applied to complex images? Obviously, these anchoring rules cannot be applied directly to complex images. A more plausible approach would be to decompose the image into components, or subimages, and then apply our rules of anchoring to each of these components. However, here we encounter several difficulties. There are a variety of kinds of components into which the image can be decomposed. We must find the appropriate one. Even if we do, it is not reasonable to expect that each component subimage can be treated in total isolation from the rest of the image. Surely, there is interaction among the subimages, and the exact nature of this interaction must be identified.

What Are the Relevant Components of a Complex Image?

Both the phenomenologist Katz (1935) and the gestalt theorists like Koffka (1935) spoke of regions of the image they called fields or frameworks. These are regions of common illumination. All the surfaces lying within a shadowed region, for example, would constitute a field. Applying the rules of anchoring within such fields of common illumination makes good intuitive sense. However, it is not immediately obvious how the visual system can identify and segregate such fields. Edge classification (Gilchrist et al., 1983) would be enormously useful here, but spelling out the rules of edge classification presents its own challenge to theory. Another approach, though related, would be to divide the image into coplanar regions, as proposed by Gilchrist (1980), but there are pitfalls here as well. What happens, for example, when a shadow falls across half of a set of coplanar regions? Intrinsic images provide yet another kind of subimage. However, each of these schemes fails when confronted with the empirical data.

Framework

We propose to define a framework in terms of the gestalt grouping principles. A framework is a group of surfaces that belong to each other, more or less. By this definition of framework, it is clear that complex images contain multiple frameworks.

These multiple frameworks can be related to each other in several ways, depending on the distribution of grouping factors in the image. Some images are divided into separate but adjacent local frameworks, like a country is divided into provinces. Some images are structured as a nested hierarchy, with several superordinate and subordinate levels. In yet other cases, the frameworks intersect one another.

Rules of Anchoring Within a Framework

In each framework, target lightness is computed according to Equation 1 or Equation 2, just as it is computed in simple images. Except by coincidence, the target will have a different computed lightness when anchored within each of these frameworks.

Local and Global Frameworks

The largest framework consists of the entire visual field and is called the *global framework.* Subordinate frameworks are called

[7] The Munsell value equivalent for a given reflectance can be obtained from a table (Judd, 1966, p. 849) or by a formula (Wyszecki & Stiles, 1967, p. 478).

local frameworks. Local frameworks are defined by local grouping factors, not by distance. There is no fixed degree of proximity within which a group of regions will be called local. A target will always be a member of at least two of these frameworks, the global framework and one or more local frameworks..

Weighting

According to our proposed model, the net lightness predicted for a given target is a weighted average of its computed lightness values in each of these frameworks, in proportion to the strength of each framework. Because the grouping factors are graded, as opposed to all or none, and because a given framework can be supported either by a single grouping factor or by several, frameworks can be said to be stronger or weaker. The strength of a framework also depends strongly on the size of the framework and on the number of distinct patches within the framework. A target that belongs to a framework containing many distinct patches will be anchored strongly by that framework. A target will be more strongly anchored by a large framework to which it belongs than to a small framework to which it belongs.

A stimulus configuration that has been frequently studied in lightness perception involves a single superordinate framework that is subdivided into two subordinate frameworks. Katz's (1911) experimental arrangement composed of adjacent lighted and shadowed fields is one example. Another is the simultaneous lightness contrast illusion consisting of two gray targets on adjacent black and white backgrounds. We can sketch out fairly well the rules of anchoring for images that contain two levels of framework, using the more convenient terms *local* and *global* even when the entire stimulus pattern does not fill the entire visual field. The following formula predicts the appearance of the target:

$$\mathrm{PR} = W_l(L_t/L_{hl} \times 90\%) + (W - 1)(L_t/L_{hg} \times 90\%), \quad (3)$$

where W_l is the weight of the local framework, $W - 1$ is the weight of the global framework, L_{hl} is the highest luminance in the local framework, and L_{hg} is the highest luminance in the global framework. When area effects apply, this formula would have to be modified as shown in Equation 2.

Belongingness and Grouping Factors

The grouping factors produce the perceptual quality of belongingness, or "appurtenance" as Koffka (1935) called it, grouping some retinal regions together and segregating others from each other. A set of coplanar surfaces appear to belong together and thus constitute a framework. A set of surfaces moving in the same direction also constitute a framework, which is based on the principle of common fate. A group of surfaces lying in shadow constitute a framework as well.

The strongest factor is probably coplanarity, at least when the luminance range is large (Gilchrist, 1980, p. 533). Classic gestalt grouping factors like proximity, good continuation, common fate, and similarity are also effective. Edge sharpness, T-junctions, and X-junctions (especially when they are ratio-invariant) are important factors in belongingness as well. Finally, many empirical results appear to require that retinal proximity be treated as a weak but inescapable grouping factor (Schirillo, Reeves, & Arend, 1990).

Importance of T-junctions. The T-junction appears to be a potent grouping factor. The general principle, in our model, is that the two occluded quadrants appear to belong to each other very strongly, whereas the occluding border seems to provide a strong segregative factor, perceptually segregating the occluding region from the two occluded quadrant regions. Todorović (1997), Ross and Pessoa (in press), and Anderson (1997) recently emphasized the role of T-junctions in such illusions, but they gave somewhat different interpretations. Todorović spoke of contrast between regions bounded by the same collinear edge, with the basis of contrast unexplained. Ross and Pessoa proposed the idea of contrast reduction at context boundaries, signaled in some cases by T-junctions. Anderson emphasized the role of T-junctions in producing scission.

Role of luminance gradients. Luminance gradients are held to segregate the luminance values at their opposite ends from each other. If two different but adjacent luminance values are divided by a sharp edge, they belong together strongly for anchoring purposes. However, if they are separated by a luminance ramp, the same two luminances will be only weakly anchored by each other.

Theoretical Value of Belongingness

There are several important theoretical advantages to the belongingness definition of a framework. First, it allows us to define frameworks in terms of retinal variables rather than in terms of a perceived variable like perceived illumination, avoiding the percept–percept coupling issue. Second, it allows a unified account of both illumination-dependent errors and background-dependent errors (simultaneous contrast). If a framework were defined as a region of common illumination, as in the usage of Katz (1935) and Koffka (1935), our anchoring analysis would not work for the simultaneous contrast display (the standard textbook version) because there both local frameworks lie in the same field of illumination.

Third, the belongingness construction allows us to bypass the problem of edge classification, even though factors like edge sharpness and coplanarity, formerly thought to underlie edge classification, now show up as grouping factors.

The anchoring model need not deny that humans can classify edges, only that lightness computation depends on edge classification. Edge classification might depend on a process that runs parallel to that of lightness computation.

The Scale Normalization Effect

We use the term *scaling* to refer to the relationship between the range of luminances in the image (or within a framework) and the corresponding range of perceived lightness values. The range of lightness values can be either expanded or compressed relative to the range of luminance values. Unless otherwise stated, our model assumes veridical, or 1:1, scaling; the range of lightness values is equal to the range of luminance values. However, we do postulate a scale normalization effect whereby the range of perceived lightness values in a framework tends to normalize on the luminance range between black and white (30:1). Whenever the luminance range within a framework is greater than 30:1, some compression occurs, but whenever the range is less than this, some expansion occurs, with the amount of compression or expansion proportional to the deviation of the stimulus range from the standard range (30:1).

Testing the Model: The Staircase Gelb Effect

Important aspects of the proposed model are illustrated by a series of experiments conducted by Cataliotti and Gilchrist (1995; Gilchrist & Cataliotti, 1994), using a display we call the *staircase Gelb effect.*[8] A contiguous series of five squares, spanning the gray scale from black to white, in roughly equal steps, was suspended in midair in a laboratory room and brightly illuminated by a homogeneous rectangular patch of light projected by an ellipsoidal theatrical spotlight mounted on the ceiling at a distance of 2.8 m from the squares, producing a luminance of 4.14 cd/m^2 at the black square. The entire lab room was illuminated by a 500-W incandescent flood lamp, but the illumination on the five squares was 30 times brighter than the ambient level at that location. The observer viewed this display with no restrictions from a distance of 4 m and indicated the lightness of each square by selecting a matching chip from a brightly illuminated Munsell chart housed in a rectangular chamber that rested on a table immediately in front of the observer.

The results are shown in Figure 6. The striking aspect in the results is the dramatic compression in the range of perceived grays. Even though the physical stimulus contains the entire gamut of gray shades from black to white, observers perceive only a range of grays from light middle gray to white.

Applying the Anchoring Model

This result makes sense if we apply our anchoring model. We can treat the stimulus squares as members of two frameworks, one local and one global. The five squares form a local group based on their proximity and their coplanarity. In addition, each is part of the global framework that includes the entire laboratory scene. The application of the model is shown schematically in Figure 7. The diagonal line (which is called the L line) shows the lightness values computed solely within the local framework, using Equation 1. The horizontal line (which is called the G line) shows the lightness values of the target squares in the global framework. They reflect the fact that, without the local group of squares, each square by itself would appear white in the global framework.[9]

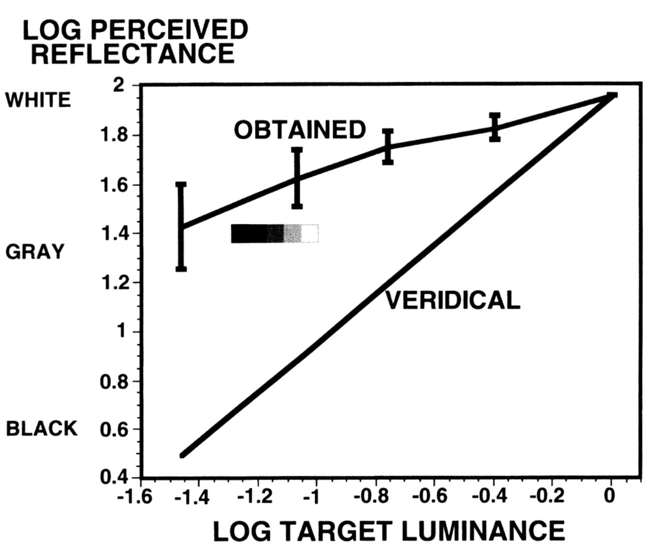

Figure 6. Perceived lightness range for five squares in spotlight is compressed relative to the actual range.

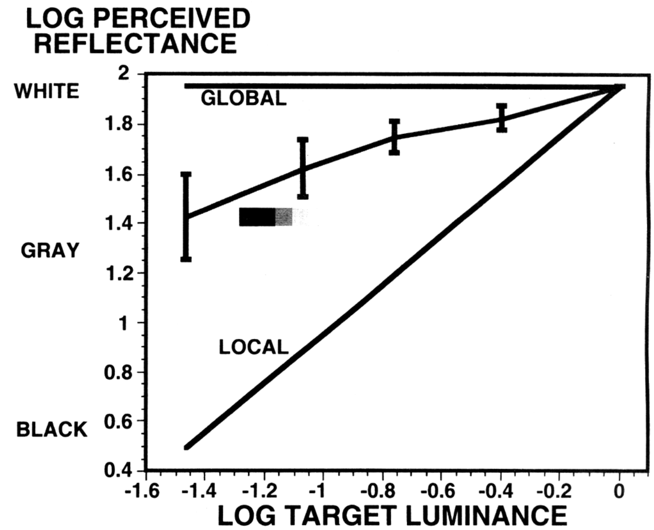

Figure 7. Schematic showing theoretical account of compression.

Compromise

Notice that the obtained value for each square lies in between its value in the local framework and its value in the global framework. This compromise lies at the heart of our new theory. In this particular case, the compromise is roughly 30% local and 70% global, but we propose that in various other situations, the balance of the compromise might shift in favor of either the local framework or the global framework, depending on the relative strength of these two frameworks.

Weighting Factors in the Staircase Gelb Effect

Gilchrist and Cataliotti (1994) conducted a series of variations on the staircase Gelb experiment to test the anchoring model by varying a series of factors that might alter the weighting of the local framework. Two of these factors, articulation and field size, were suggested by the early lightness perception literature, especially in the work of Katz (1935). Two other factors, which we call

[8] We thank William Ross and Luiz Pessoa for suggesting this name.

[9] The horizontal slope of the global line deserves comment. If only the highest luminance rule of anchoring were applied, the global line would have to have at least some slope because the white square is the highest luminance in the global framework and each of the other four squares is darker than the white. However, we presume that relative area plays an anchoring role in the global framework also. Because the five squares take up only a small portion of the global area, we assume that they exert very little effect on the global anchor. Thus, all five squares are approximately at or above the global anchor, and it seems reasonable to present the global line as horizontal. It may be that the G line should have a slight slope, but we defer this issue.

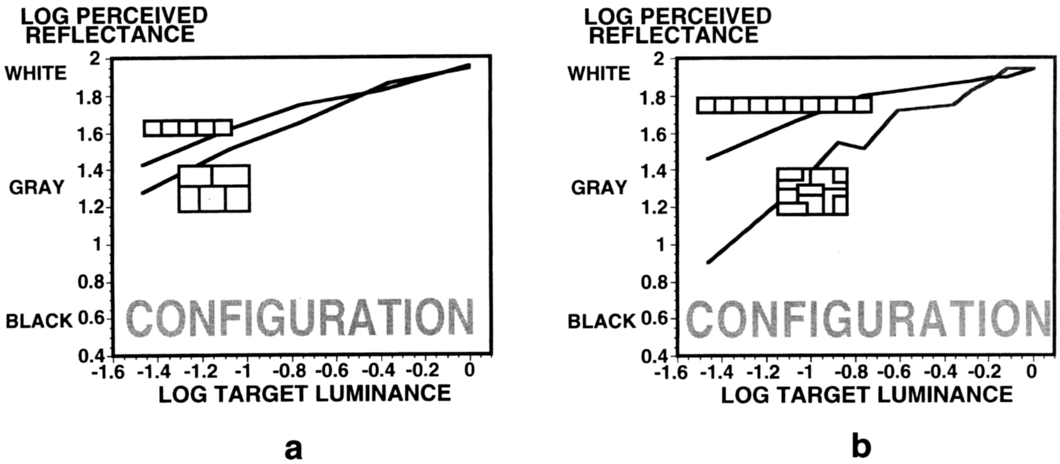

Figure 8. Dependence of compression on configuration of targets. Mondrian configuration produces stronger local anchoring.

configuration and *insulation,* were uncovered in the course of the investigation.

Configuration. Apparently, the squares constitute a stronger framework if they are arranged in a Mondrian pattern than if they are simply arranged in a line. This can be seen in Figure 8. Notice that both in the case of 5 squares and in the case of 10 squares, the obtained data fall closer to the L line than to the G line, suggesting that the local framework is stronger with a Mondrian configuration. Several additional experiments are needed to sort out whether the Mondrian configuration produces better constancy than the linear configuration because of the greater number of adjacent ratios, because of the scrambling of the luminance staircase, or because of some other factor.

Articulation. Three Mondrians—containing 2, 5, and 10 target surfaces, respectively—were presented under the same basic conditions. The results, shown in Figure 9, make it clear that the strength of the local framework depends strongly on the number of squares in the group, even when the luminance range within the group is held constant. Our data indicate that the crucial factor is the number of different surfaces, not the number of different gray levels, but this result needs to be confirmed under a wider range of conditions. We call this factor articulation, defined simply as the number of different surfaces in a framework, in deference to Katz (1935).

Katz (1911) argued that the greater the articulation within an illumination frame of reference, the higher the degree of lightness constancy. We are modifying Katz's usage of this concept somewhat. Our proposal is that the higher the degree of articulation in a target's framework, the more the lightness of the target is anchored solely within that framework.

Field size. According to Katz (1935), the degree of lightness constancy within a given field of illumination depends on the size of the field. However, as Rock (1975, 1983; Rock & Brosgole, 1964) has shown so often for various factors in perception, size can be defined either in retinal terms or in phenomenal terms. Katz believed that both of these meanings of size are effective in lightness constancy, and hence he offered his two laws of field size. The first law holds that the degree of constancy varies with the retinal size of a field. He supported it with the observation that if one looks through a neutral density filter held at arms length and then slowly brings the filter toward the eye, the degree of lightness constancy for surfaces seen within the boundaries of the filter increases as the filter comes to occupy a larger proportion of the visual field.

The second law holds that constancy varies with perceived size. This he demonstrated by keeping the filter at arms length while he slowly walked backward away from a wall containing various surfaces. As the perceived size of the region of the wall seen through the filter increases, so does constancy, even though retinal size is held constant.

Although his second demonstration established the effectiveness of perceived size with retinal size held constant, his first demonstration did not establish retinal size with perceived size held

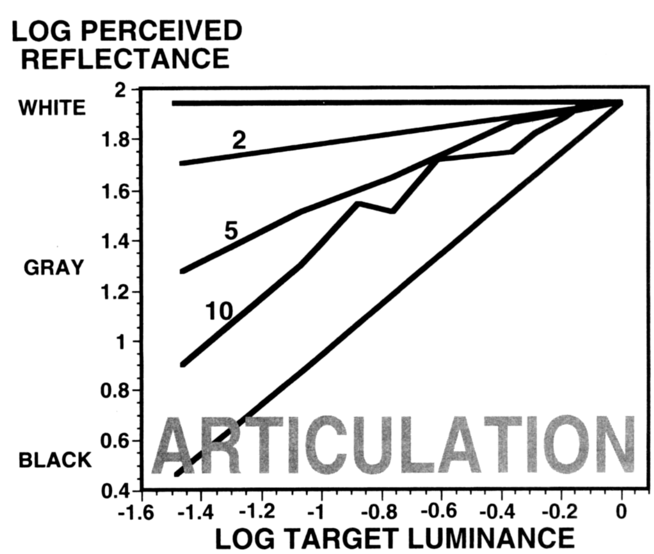

Figure 9. Dependence of compression on number of squares. Greater articulation (more squares) produces stronger local anchoring.

constant. When the filter is drawn closer to the observer's eye, the total area of the surfaces seen through the filter grows both in retinal and perceived size.

Perceived size, not retinal size, is crucial. In several experiments, we have found perceived size, not retinal size, to be effective (when perceived size is controlled). We repeated the experiment with the five-square Mondrian but used a Mondrian five times larger both in width and height. Results from the small Mondrian and the large Mondrian are shown in Figure 10 (Gilchrist & Cataliotti, 1994). The increase in size produced a significant darkening for only the black square, but there appears to be a trend for the other squares as well. We obtained little or no difference when we increased size simply by moving the observer closer to the Mondrian, which increased the net size of the display with little or no change in its perceived size.

Bonato and Gilchrist (1999) found higher luminosity thresholds for targets of larger perceived area but not for targets of larger retinal area. Bonato and Cataliotti (in press) showed the importance of phenomenal size in a different way. They found a higher luminosity threshold for a region perceived as ground than for a region perceived as figure. Although the area of the two regions was the same in the display (shown in Figure 11), the phenomenal area of the ground region is greater because it is perceived to extend behind the figural (face) region. This hidden part of the background can be called its amodal area (Kanizsa, 1979). A quantitative analysis of the Bonato and Cataliotti results indicates that the functional area of the ground region (as this bears on the luminosity threshold) includes only some of the area behind the figure, not all of it. This finding is consistent with that of Shimojo and Nakayama (1990), who used an apparent motion display to determine how far a ground region is perceived (functionally) to extend behind a figure–ground contour.[10]

Insulation. Gilchrist and Cataliotti (1994) discovered that a white border surrounding the group of squares seems to insulate it from the influence of the global framework. A border of black paper has no such effect, as can be seen in Figure 12. The same insulation effect applies to the single square used in the basic Gelb effect as well. Cataliotti and Gilchrist (1995) found that when a white square is placed next to the black Gelb square, the darkening of the black square is less than half what would be commensurate with the luminance ratio between them. However, both they and McCann and Savoy (1991) found that when the white region completely surrounds the Gelb square, the darkening effect is highly commensurate with the luminance ratio between them.

At this point, this insulation effect remains little more than an empirical result; we have no deeper explanation for it. But the effect is apparently not reducible to local contrast. When the five squares with a white border were compared with a window panes arrangement in which contact between each of the darker targets and white was maximized and with a concentric arrangement in which the contact was minimized, no differences were found, as can be seen in Figure 13.[11] The key requirement seems to be merely that the inner framework is completely enclosed by a white border.

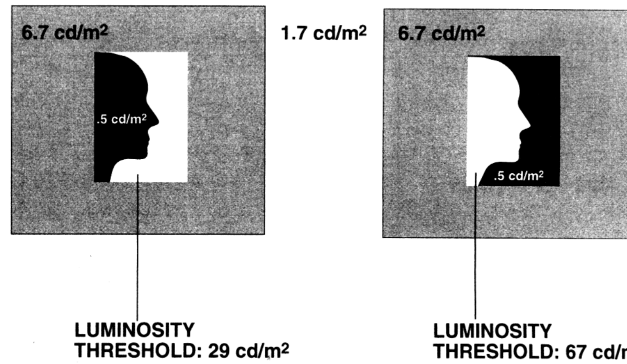

Figure 11. Greater perceived area produces higher luminosity threshold. From "The Effects of Figure/Ground, Perceived Area and Target Saliency on the Luminosity Threshold," by F. Bonato and J. Cataliotti, in press, *Perception & Psychophysics.*

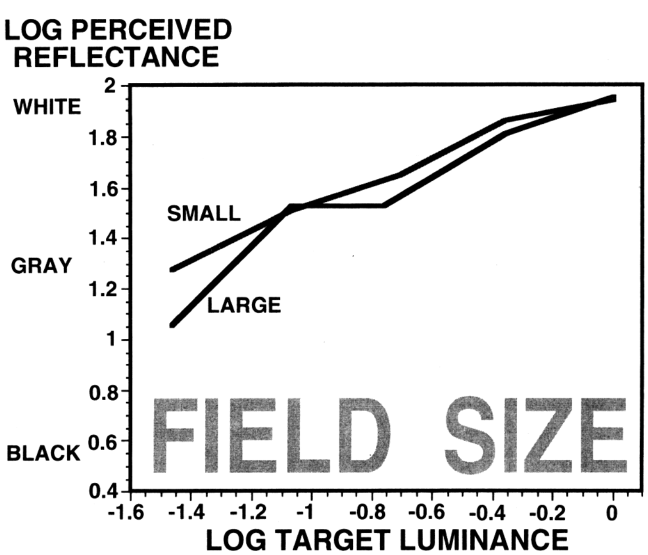

Figure 10. Dependence of compression on size of display. Larger field size produces stronger local anchoring.

Anchoring and the Pattern of Lightness Errors

We began by noting two weaknesses of the intrinsic image models: the anchoring problem and the errors problem. We have quite an ironic situation here. The anchoring model evolved out of an attempt to fill an important gap in the intrinsic image models, namely, the lack of anchoring rules. Yet, the resulting anchoring

[10] This revised definition of area must, of course, be applied also to our domes experiments (see Figure 3, especially d and f). But two things need to be kept in mind. First, only part of the amodal area is attributed, functionally, to the background region. Second, there is a certain degree of figure–ground ambiguity in the displays that involved the large oval (Figures 3d and 3f). It must be kept in mind that these figures are only schematic, and in the actual dome, the large oval covered most of the visual field.

[11] There does appear to be a weak trend in this figure. However, in a separate replication of this experiment, no such trend appeared.

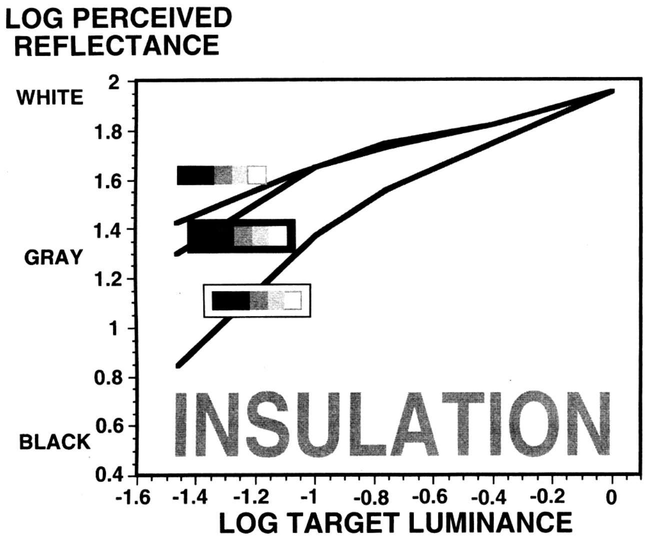

Figure 12. Insulation: white border produces stronger local anchoring; black border has no effect.

model appears to be inconsistent with an intrinsic image model. If our investigation of the anchoring problem has undermined the intrinsic image models, a study of the errors problem promises to be no more accommodating, given that the intrinsic image models are, at base, models of veridical perception.

The errors problem is the challenge of explaining the empirical pattern of errors in lightness perception, under a wide range of stimulus conditions. What kind of visual processing could achieve the impressive degree of lightness constancy shown by human observers and at the same time produce just that pattern of errors they show?

We are aware that the very concept of perceptual error is fraught with philosophical contention. Though we do not address this debate here, we give a very concrete definition of a lightness error: the difference between the actual reflectance of a target surface and the reflectance of the chip selected by the observer from a well-lighted Munsell chart as having the same apparent lightness as the target.

One cannot find in the lightness literature either a survey of the pattern of errors that have been obtained in empirical work or a survey of the pattern of errors predicted by the main theories of lightness perception. A comparison of predicted and obtained errors shows a profound mismatch; many of the important errors predicted by theories simply do not occur, and much of the pattern of obtained errors has remained untouched by lightness theory.

Given the vast range of empirical findings on lightness errors, we focus on achieving a single explanation for two broad classes, illumination-dependent errors and background-dependent errors, represented respectively by the Katz (1911) experimental arrangement and the textbook version of simultaneous lightness contrast. Although contrast theories have claimed to account for both of these, they can account, at best, only for the direction, not the magnitude.

Gilchrist (1988): Failures Due to a Classification Problem

Gilchrist (1988) proposed a common explanation for illumination-dependent errors and background-dependent errors that was based on edge classification. This approach is problematic because edge classification, at least for coplanar edges, is all-or-none, whereas constancy failures are graded. The proposal involves a graded concept of edge classification. It was based on an experiment in which targets were placed on adjacent bright and dark backgrounds that appeared to differ in reflectance in one condition but illuminance in another condition. Thus, the two conditions allowed for a test of illumination-independent constancy and background-independent constancy under identical stimulus conditions, save for changes in the larger context surrounding the two backgrounds. Veridical performance would have produced ratio matching for illumination-independent constancy but luminance matching for background-independent constancy. The actual results fell between these two poles in both cases, but there was a further unique feature of the data. The deviation from ratio matching in the illumination-dependent constancy condition was exactly equal to the deviation from luminance matching in the background-dependent situation, as shown in Figure 14.

Gilchrist (1988) noted that this pattern in the data could be the result of a compromise between two competing processes, one process that produces ratio matches and one that produces luminance matches. The first of these would be invoked if the edge between the backgrounds were classified as an illuminance edge, the second if it were classified as a reflectance edge. The obtained data are consistent with an incomplete classification in both cases. The deviation from ratio matching would be predicted if the illuminance border dividing adjacent illuminated and shadowed fields were not classified completely as an illuminance border but were classified to some small degree as a reflectance border. Likewise, the deviation from luminance matching would also be

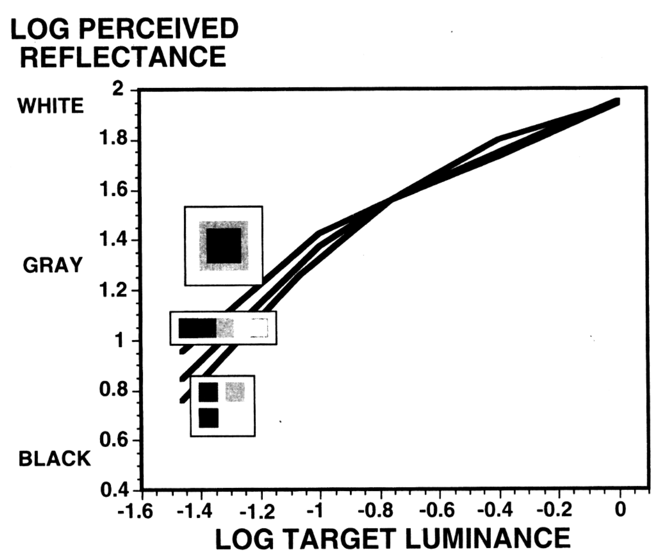

Figure 13. Test of contrast interpretation of borders effect. Lightness of black target shows little or no dependence on degree of contact with white border.