1. Introduction: Theoretical Developments; From Unconscious Inference to Intrinsic Image:

- 1.1. Basic ambiguity: luminance vs. lightness; lightness constancy
- 1.2. Inferring the illuminance: Helmholtz
- 1.3. Hering's paradox
- 1.4. Gestalt and the relational approach
- 1.5. Contrast theories versus relational theories
- 1.6. Intrinsic image theories
    - 1.6.1. Edge integration.
    - 1.6.2. Edge classification.
    - 1.6.3. Parsing into layers
- 1.8. Two weaknesses of the intrinsic image models.
    - 1.8.1. The problem of errors.
    - 1.8.2. The missing anchor.

2. The Anchoring Problem

- 2.1. A concrete example
- 2.2. Mapping luminance onto lightness
- 2.3. Wallach on anchoring: Highest Luminance Rule
- 2.4. Helson on anchoring: Average Luminance Rule
- 2.5. Anchoring in intrinsic image models
    - 2.5.1. Anchoring applied to reflectance layer.
- 2.6. Absolute lightness versus absolute luminance.
- 2.7. Anchoring versus scaling

3. The Rules of Anchoring in Simple Displays

- 3.1. What are minimal conditions?
- 3.2. Highest Luminance versus Average Luminance
- 3.3. Luminosity Problem: Direct Contradiction to HLR
    - 3.3.1. Wallach on increments: HLR does not apply
    - 3.3.2. Common factor for decrements and increments: Geometric, not photometric

4. Area effects

4.1. Surround rule fails

4.2. Highest luminance plus area

4.3. The Area Rule

4.3.1. Other empirical findings

4.4. Luminosity and the Area Rule

4.4.1. Area Rule implies luminosity:

4.4.2. Luminosity induction versus grayness induction.

4.4.3. Phenomena related to the Area Rule

4.4.4. Luminosity threshold and area: Bonato and Gilchrist

4.4.5. Maximum luminance is not the same as the anchor

4.5. Anchoring rules for simple images: A summary

5. Anchoring in Complex Images: A New Theory

5.1. What are the relevant components of a complex image?

5.2. Framework

5.3. Local and global frameworks

5.4. Weighting

5.5. Belongingness and grouping factors

5.5.1. Importance of T-junctions

5.5.1. Role of luminance gradients

5.6. Theoretical value of belongingness.

5.7. The scale normalization effect.

6. Testing the model: The staircase Gelb effect

6.1. Applying the anchoring model

6.2. Compromise

6.3. Weighting factors in the staircase Gelb effect

6.3.1. Configuration.

6.3.2. Articulation.

6.3.3. Field Size.

6.3.4. Perceived size is crucial, not retinal size

6.3.5. Insulation

7. Anchoring and the Pattern of Lightness Errors

7.1. Gilchrist (1988): Failures due to classification problem

7.1.2. Staircase Gelb data inconsistent

7.1.3. Anchoring model consistent with Gilchrist (1988) data.

7.2. Source of Errors

8. The Model Applied to Illumination-Dependent Errors

8.1. Errors associated with level of illumination

8.2. Errors associated with reflectance of target.

8.3. Errors associated with reflectance of backgrounds

8.3.1. The key: Increments versus decrements

8.4. Errors associated with articulation level

8.4.1. Burzlaff experiments.

8.4.2. Arend & Goldstein experiments.

8.4.3. Schirillo experiments.

8.4.4. Application of model to depth and lightness.

8.5. Errors associated with field size

9. The Model Applied to Background-Dependent Errors (Simultaneous Lightness Contrast)

9.1. Anchoring component

9.2. The scale normalization component

9.3. Predictions of the model

9.3.1. The locus of the error.

9.3.2. The size of the error.

9.3.3. No contrast effect with double increments

9.3.4. Segregating effect of luminance ramps

9.3.5. Belongingness creates the illusion

9.4. The Benary effect

9.5. White’s illusion.

9.6. Checkerboard contrast.

9.7. Agostini and Proffitt: Common fate.

9.8. Laurinen and Olzak.

9.9. Wolff illusion

9.10. Depth manipulations:

9.11. Adelson’s corrugated Mondrian

9.11.1. Wishart experiments.

10. Brightness Induction: Contrast or Anchoring?

10.1. Induction

10.2. Area effects

10.3. Separation

10.4. Depth separation: Gogel and Mershon

10.5. Gelb effect: Contrast versus anchoring

10.6. Summary: application of model to reduction conditions.

10.6.1. Lightness versus brightness

11. Tests of scaling normalization effect

11.1. Crispening effect

12. Problems for the Intrinsic Image Models

12.1. Problem of the staircase Gelb effect data.

12.2. Problem of area effects.

12.3. Problem of articulation effects.

12.4. Problem of errors.

12.5. A response paradox.

13. New Model vs. Intrinsic Image Models

13.1. Classified edge integration

13.2. A Critical Test

13.2.1. Method and results: New model wins.

13.2.2. Low articulation replication.

13.2.3. Review of 1983 data.

14. Evaluating The Model

14.1. Range of application.

14.2. Prediction of errors.

14.3. Unification of constancy failures.

14.4. Rigor and concreteness.

14.5. The role of perceived illumination.

14.6. Choice of scaling rule for the global framework.

15. Debt to the Early Literature

15.1. Koffka

15.2. Compromise and intelligence

15.3. The aperture problem for lightness

An Anchoring Theory of Lightness Perception

[1]Gilchrist, A., [2]Kossyfidis, C., [3]Bonato, F., [4]Agostini, T., [3]Cataliotti, J.,

[1]Li, X., [5]Spehar, B., [1]Szura, J.

[1]Rutgers University

[2]Deceased

[3]St. Peters College

[4]University of Trieste

[5]University of New South Wales

This paper is dedicated to the memory of Christos Kossyfidis.

Running Head: AN ANCHORING THEORY OF LIGHTNESS PERCEPTION

Abstract

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 two 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.

1. Introduction: Theoretical Developments; From Unconscious Inference to Intrinsic Image:

1.1. Basic ambiguity: luminance vs. lightness; lightness constancy

We present here 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.

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. But 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, we 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.

1.2. 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. But as an account of human lightness, the theory has been fraught with both logical and empirical difficulties.

1.3. Hering’s paradox

Hering (1874) argued that Helmholtz’ position is illogical because, given the luminance at any point, one would need to know the surface reflectance to infer the illumination. But one would also need to know the illumination to infer the surface reflectance. Hering emphasized the role of sensory mechanisms like pupil size and adaptation. But he emphasized lateral inhibition as the main mechanism, many years before its actual discovery, thus creating the prototype of the modern contrast theories.

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, especially in relation to his historical context. 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 mechanisms plus memory color are ruled out. Katz’s theoretical perspective, though initially closest to that of Helmholtz, became strongly influenced by the Gestalt theorists.

1.4. 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. As Marr (1982, p. 259) has noted: "It is a widespread and time-honored view, going back at least to Ernst Mach, that object color depends upon the ratios of light reflected from the various parts of the visual field rather than on the absolute amount of light reflected." But Koffka (1935, p. 244) was the first to base a theory of lightness perception strictly on relative luminance: "Our theory of whiteness constancy will be based on this characteristic of colours...that perceived qualities depend upon stimulus gradients." 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 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 the first edition 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 1940's, 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 paper, 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.

1.5. Contrast theories versus relational theories

Hering (1874) is often described as one of the earliest to point out the importance of relative luminance. But, as Koffka (1935, p. 245) wrote about Hering: "contrast...implies an explanation not in terms of gradient, but in terms of absolute amounts of light." This is explicit in more recent contrast theories within the Hering tradition. Cornsweet (1970, p. 303) defines the correlate of perceived lightness as "the frequency of firing of the spatially corresponding part of the visual system (after inhibition)." Likewise Jameson and Hurvich (1964) seem to acknowledge a fundamental role for relative luminance but their position is that the lightness values produced by a given luminance ratio in the stimulus depends, in the end, on the absolute luminance values (Jameson & Hurvich, 1961). This was the central point of their celebrated but essentially unreplicable (Flock & Noguchi, 1970; Gilchrist & Jacobsen, 1988; Haimson, 1974; Heinemann, 1988; Jacobsen & Gilchrist, 1988a and 1988b; Noguchi & Masuda, 1971) report.

Relational theories like those of Koffka, Helson, and Wallach can be driven equally well by an input consisting either of absolute or relative luminance values. Contrast theories, while sometimes couched in relational terms, require absolute luminance information as well.

Although lightness perception continues to be attributed to contrast mechanisms by non-specialists, those who study surface lightness have long recognized that lateral inhibition, although crucial to the encoding of stimulus values, plays no more substantial role.

1.6. 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). And yet there has been a reluctance to embrace such a simple idea.

In the past two 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 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. Ross and Pessoa (personal communication) 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).

1.6.1. Edge integration.

Wallach's ratio rule deals effectively with adjacent luminances. But 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 papers, 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.

1.6.2. Edge classification.

Gilchrist (1977, 1979, 1980) demonstrated empirically an observation that had been made by Koffka in 1935 (p. 248) that "not all gradients are equally effective as regards the appearance of a particular field part....," and, in another passage (p. 246): "Clearly two parts at the same apparent distance will, ceteris paribus, belong more closely together than field parts organized in different planes." 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 observation that some luminance ratios are relatively effective in determining surface lightness while 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 while 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.

1.6.3. 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 (1990) have recently proposed an elegant scheme that bears a striking resemblance to Bergström's model. They liken 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 the cost of these three methods is not the same, and this sets 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, Gilchrist, Barrow and Tenenbaum, and Adelson and Pentland 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, et al, 1985) in his analysis of perceived transparency. We will refer to these models as intrinsic image models. An excellent discussion of them can be found in a recent chapter by Arend (1994).

1.8. 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: (1) they cannot account for errors in lightness perception and (2) they have no anchoring rule. We consider these in turn.

1.8.1. 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 byproduct 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, to that extent 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. This pattern is 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. But no such prospect is in sight. For example there is no coherent approach that can explain both illumination-dependent failures of constancy (Type 1), and background-dependent failures of constancy (Type 2), 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.

1.8.2. 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 will bring us back to the errors problem.

We will 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 empirical results.

2. The Anchoring Problem

2.1. 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 the 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 five to one ratio (see Figure 1). This five to one 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 five to one ratio. For example, if the five represents white then the one represents middle gray. But the five might represent middle gray, in which case the one will represent black. Indeed it is even possible that the one represents white and the five 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.

2.2. 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. 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 2. Consider the following remark by Koffka (1935, p. 255) in relation to Figure 2: "...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." Anchoring concerns where this sliding relationship comes to rest. Koffka suggests here 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.

2.3. 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) have also adopted this rule, which we will call the Highest Luminance Rule.

2.4. Helson on anchoring: Average Luminance Rule

Although Helson 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 (Bruno, 1992; Bruno, Bernardis, & Schirillo, in press; Brown, 1994; Schirillo & Shevell, 1996). Land (1983) reverted to the average luminance rule in a later version of Retinex theory.

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.

2.5. Anchoring in intrinsic image models

2.5.1. Anchoring applied to reflectance layer.

None of the intrinsic image models contains 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.

2.6. Absolute lightness versus absolute luminance.

The issue of relative versus absolute lightness, central to the anchoring problem, is independent of the issue of whether relative or absolute luminance values are encoded at input. The former applies to the output, the latter applies to the input. Normalization typically refers to the fact that absolute luminance values are lost in the encoding of retinal luminance contrasts, whereas the anchoring problem concerns lightness values, not luminance values. It should be noted that even if absolute luminance values were encoded at input (a dubious assumption; see Shapley and Enroth-Cugell, 1984; Whittle & Challands, 1969) or somehow recovered later (perhaps by combining relative luminance information with adaptive state information), the anchoring problem would remain completely unsolved, due to the lack of correlation between luminance and lightness.

2.7. 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, that we call the scaling problem, is illustrated in Figure 3. 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. Re-scaling 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, use the term gamut expansion), which is just the reverse.

While 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 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.

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 and 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, unlike Wallach, 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.

3. The Rules of Anchoring in Simple Displays

Our approach will be 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 will find that, for simple images, anchoring depends on two dimensions of the stimulus: relative luminance and relative area.

3.1. 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 non-zero luminance. Wallach (1963 p. 112) 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..." 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 (1971, p. 146) has expressed a common view that the "simplest experimental arrangement for studying induction effects" 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.

3.2. Highest Luminance versus Average Luminance

Which rule is correct under such minimal conditions, Wallach's highest-luminance-as-white rule, or Helson's average-luminance-as-middle-gray rule? This question cannot be answered when the stimulus contains the entire gamut of gray shades from white to black because in that case the two rules make the same predictions. The most direct test of these two rules is to present an observer with a truncated gray scale; an array of luminances whose range is substantially less than the 30:1 range between white and black.

Combining a truncated gray scale with simple conditions, we placed observer's heads inside a large acrylic hemisphere, the inside surface of which was divided into two halves of equal size (Li and Gilchrist, in press). One half was painted completely matte black; the other half middle gray. The middle gray half was seen as fully white and the black half as a darkish middle gray. 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 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.

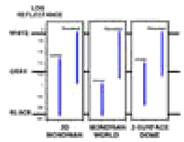

The two rules have been tested indirectly in several experiments on "equivalent surrounds." Bruno (1992) asked observers to make a brightness match between two target squares of equal luminance, one surrounded by a homogeneous region and one embedded in a checkerboard or Mondrian-like region. The luminance of the homogeneous surround that is chosen by the observers as having the same effect on a target as the checkerboard is a luminance value close to that of the highest luminance in the checkerboard not the average. 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 is increased from 0% to 100%. Their results for increments (regions with a luminance higher than that of the surround) agree with those of Bruno, but their results for decrements show a different pattern (see also Bruno, Bernardis, and Schirillo, in press).

McCann (1994) has recently shown 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. But 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).

3.3. Luminosity Problem: Direct Contradiction to HLR

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. But the fact is that we perceive various regions within the visual field to be self-luminous. This fact represents a serious challenge to the highest luminance rule.

3.3.1. Wallach on increments: HLR 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. But Wallach reported informally that when the disk is an increment, that is brighter than the surrounding annulus, it appears luminous. The implication is that the highest luminance rule applies to decrements but not to increments. Here are his words: "We have seen that when two areas of different stimulus intensity exert their influence on each other, the area of lower intensity will be gray and the area of high intensity will be white. However, the correspondence between white and gray goes no further. When the difference in intensity is varied, the value of the gray changes; the gray may even turn black. If the spot is more intense, it will show no color other than white. An increase in the intensity difference will cause a different kind of quality, that is, luminosity, to appear in addition to the white and a further increase will merely cause luminosity to become stronger." (Wallach, 1976, p. 8) Thus, the luminosity that occurs with increments directly contradicts Wallach’s Highest Luminance Rule.

3.3.2. 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 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 rule.

Wallach’s 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.

4. Area effects

In the disk/ganzfeld experiment, the figure/ground distinction was confounded with relative area. Li and Gilchrist (in press) conducted a series of experiments to separate relative area from figure/ground. To create a stimulus pattern that fills the entire visual field, the observer’s head was placed inside a large acrylic hemisphere. A simple two-part pattern was painted on its opaque interior. Relative area was varied while figure/ground arrangements were held constant. Schematic representations of these stimuli are shown in Figure 5, along with the results, given in Munsell values.

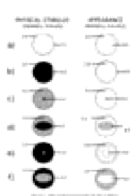

4.1. Surround rule fails

The net result is that the surround rule was completely undermined. In the incremental large oval condition, illustrated in Figure 5f, 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 is seen clearly in the split dome conditions, shown in Figure 5a and 5b. As the darker region increases its area from half the visual field (Figure 5a) to most of the visual field (Figure 5b), its lightness moves from Munsell 4.5/ to 7.8/. This finding brings to mind Helson’s (1964, p. 292) 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."

4.2. Highest luminance plus area

The highest luminance rule was rescued but combined with a tendency for the largest area to appear white. Thus anchoring under minimal conditions (two regions in the visual field) appears to depend upon both photometric and geometric factors. We can say that the larger a surface, the lighter it appears, or "the larger the lighter" for short. But "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.

4.3. 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. At the same time the lighter region first appears white, then a fluorent white and finally, self-luminous.*

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.

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

4.3.1. Other empirical findings

There are at least ten other papers in the literature concerning the influence of area on lightness or brightness. We will 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. But 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, 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. But 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 (Heinemann, 1971; Freeman, 1967, p. 173).

We will 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 will be considered in Section 10 as part of a general review of that literature.

4.4. Luminosity and the Area Rule

4.4.1. 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 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.

4.4.2. Luminosity induction versus grayness 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 might be called luminosity induction; it occurred in Gilchrist and Bonato’s disk/ganzfeld experiments when the disk was an increment. Alternatively, the lighter region might remain perceptually constant while the darker region becomes darker and darker gray. This occurs in disk/annulus experiments (Heinemann, 1955; Wallach, 1948) when the disk is a decrement. Heinemann calls this brightness induction, but we will use the term grayness induction to distinguish it from luminosity induction, and to capture the notion of two opposing directions, grayness and luminosity.

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 luminosity induction in the lighter region, with only a small amount of grayness induction. But when the area of the darker region is small relative to the lighter region, there is very little luminosity induction; most of the effect is expressed as grayness induction in the darker region.

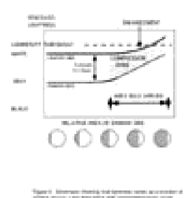

Figure 6 represents what is currently our best understanding of the appearance of the two regions in a simple framework as the relative area shifts from the lighter region to the darker. Typically a figure of this kind would plot perceived lightness and luminosity as a function of luminance ratio, with 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. Now, 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.

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. But 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. 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.

4.4.3. 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 pre-luminous 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) have recently reported a phenomenon that they describe as a novel compression mechanism in the luminance-brightness mapping. They call it brightness indention. Using a display that consisted of two disks within a ganzfeld, Schouten and Blommaert found that when the disks are both 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.

4.4.4. 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 and Gilchrist, in press) 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.

4.4.5. 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.

It should be further noted that the anchor, or luminance value that appears white, need not appear in a framework as an actual surface. For example, if a small white disk is placed in the center of a dome painted black, the disk will appear luminous and the black dome will appear light gray. Perceived white lies between these two perceived levels but is not represented by a physical surface. Within a framework, the lightness of a target surface is given by the ratio between the target surface luminance and the luminance of the anchor whether or not the anchor has the highest luminance and whether or not the anchor luminance is actually represented by a surface.

4.5. 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 that, according to the formula:

$$PR = Lt/Lh \times 90\% \quad (1)$$

where PR is perceived reflectance, Lt is the luminance of the target, Lh is the highest luminance in the framework, and 90% is the reflectance of white.

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 6, can be summarized algebraically as follows:

$$PR = (100-Ad)/50 * (Lt/Lh \times 90\%) + (Ad-50)/50 * (90\%) \quad (2)$$

where Ad is the area of the darker region, as a percentage of the total area in the field. The formula simply says that if Ad is 50% of the total area, the perceived reflectance of the darker region is just as it is given in Formula (1) above. As Ad 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 Ad grows, the lighter region takes on additional qualities, first fluorence, and finally, self-luminosity. This qualitative change, not surprisingly, is difficult to capture mathematically.

5. 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. But 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 sub-images, and then apply our rules of anchoring to each of these components. But 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 sub-image can be treated in total isolation from the rest of the image. Surely there is interaction among the sub-images and the exact nature of this interaction must be identified.

5.1. 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. But 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 sub-image. As we shall see, however, each of these schemes fails when confronted with the empirical data.

5.2. 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 (see Figures 25 and 28).

5.3. Local and global frameworks

The largest framework consists of the entire visual field and will be called the global framework. Subordinate frameworks will be called 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. In each framework, target lightness is computed according to Formula 1 or Formula 2 (Section 4.5), 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.

5.4. 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 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 such 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:

$$PR = Wl(Lt/Lhl * 90\%) + (W-1)(Lt/Lhg * 90\%) \quad (3)$$

where Wl is the weight of the local framework, W-1 is the weight of the global framework, Lhl is the highest luminance in the local framework, and Lhg is the highest luminance in the global framework. When area effects apply, this formula would have to be modified as shown in Formula 2 (Section 4.5).

5.5. Belongingness and grouping factors

The grouping factors produce the perceptual quality of belongingness, or appurtenance as Koffka (1935) called it. 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, 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, as we shall see. Finally, many empirical results appear to require that retinal proximity be treated as a weak but inescapable grouping factor. Grouping factors can segregate as well as integrate.

5.5.1. Importance of T-junctions

The T-junction appears to be a potent grouping factor. In our model, T-junctions function as illustrated in Figure 7. The general principle seems to be that the two "occluded" quadrants (B & C) appear to belong to each other very strongly while the "occluding" border seems to provide a strong segregative factor, perceptually segregating the occluding region (A) from the two occluded quadrant regions. Todorovic' (1997), Ross and Pessoa (in press), and Anderson (1997) have recently emphasized the role of T-junctions in such illusions, but they have given somewhat different interpretations. Todorovic' speaks of contrast between regions bounded by the same collinear edge, with the basis of contrast unexplained. Ross and Pessoa propose the idea of contrast reduction at context boundaries, signaled in some cases by T-junctions. Anderson emphasizes the role of T-junctions in producing scission.

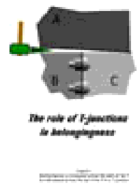

5.5.1. Role of luminance gradients

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. But if they are separated by a luminance ramp, the same two luminances will be only weakly anchored by each other.

This may seem backwards. One might argue that when

5.6. 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 bypasses the problem of edge classification. At the same time grouping factors must be identified that create the visual experience of a special region of illumination such as a shadows. The segregating role of the penumbra (luminance gradient) is one. Another might be called luminance similarity. All regions in the shadow share a lowered luminance relative to regions outside the shadow.

It is true that we now have a lot of evidence, both empirical and phenomenological, that edges are perceptually classified. And yet, the basis of edge classification has never been fully spelled out. Indeed, our emphasis on belongingness may well provide a new angle from which to attack the classification problem, touching, as it does, on the central problem of perceptual structure. Moreover, the anchoring model need not deny that humans can classify edges. Rather the model carries the more modest implication that lightness computation does not depend on edge classification. Edge classification might depend on a process that runs in parallel to that of lightness computation.

5.7. The scale normalization effect.

The term scaling refers 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).

6. 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 will call the staircase Gelb effect. 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. 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 entire lab room was normally illuminated by fluorescent lighting but the illumination on the five squares was thirty times brighter than the ambient level.

The results are shown in Figure 8. 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.

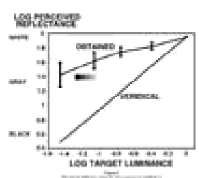

6.1. Applying the anchoring model

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, their coplanarity, and most likely their similar luminance values. In addition each is part of the global framework that includes the entire visual field. The application of the model is shown schematically in Figure 9. The diagonal line (which will be called the L-line) shows the lightness values computed solely within the local framework, using formula (1) from Section 4.5. The horizontal line (which will be 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.

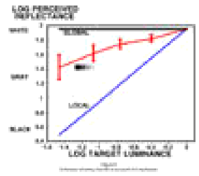

6.2. Compromise

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 upon the relative strength of these two frameworks.

6.3. 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 David Katz (1935). Two other factors, that we call configuration and insulation, were uncovered in the course of our investigation.

6.3.1. 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 10. Notice that both in the case of five squares and in the case of ten 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 further 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.

6.3.2. Articulation.

Three Mondrians, containing 2, 5, and 10 target surfaces respectively, were presented under the same basic conditions. The results, shown in Figure 11, 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 number of different surfaces, not number of different gray levels, but this 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). Perhaps this usage fails to capture the richness of Katz's concept of articulation, but our goal is to be more operational than Katz while using his term in order to recognize his contribution.

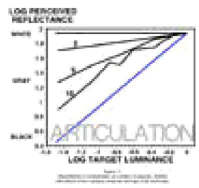

Katz argued that the greater the articulation within an illumination frame of reference, the higher the degree of lightness constancy. We are modifying Katz's usage of this concept somewhat. Our proposal is that the higher the degree of articulation in a target's framework, the more the lightness of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

6.3.3. Field Size.

According to Katz (1935), the degree of lightness constancy within a given field of illumination depends on the size of the field. But as Rock (1975, 1983; Rock & Brosgole, 1964) has shown so often for various factors in perception, size can be defined in either 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 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 though 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 demonstrates by keeping the filter at arms length while he slowly walks backward away from a wall containing various surfaces. As the perceived size of the region of wall seen through the filter increases, so does constancy, even though retinal size is now held constant.

Although his second demonstration establishes the effectiveness of perceived size with retinal size held constant, his first demonstration does not establish retinal size with perceived size held 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.

6.3.4. Perceived size is crucial, not retinal size

In several experiments we have found perceived size to be effective but not retinal size (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 12 (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. Of course, if the observer were moved so close that the five squares display fills the entire visual field, this would constitute a qualitatively different stimulus and perceived lightness values would change substantially.

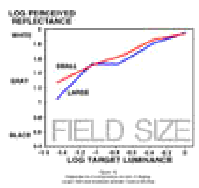

When Bonato and Gilchrist (1994) found higher luminosity thresholds for targets of larger area, their targets were larger both in visual angle as well as perceived size. In a follow-up experiment (Bonato & Gilchrist, in press) they tested luminosity thresholds for the disk in a disk/annulus display. They varied, in a controlled way, both the retinal size of the display and its perceived size, finding the luminosity threshold to vary with perceived size but not retinal size.

Bonato and Cataliotti (submitted) 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 13), 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 those of Shimojo & 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.

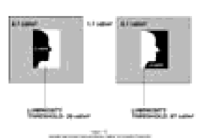

6.3.5. Insulation

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 14. Thus, when a region of maximum luminance completely surrounds a group of specially lighted surfaces, it seems to dramatically reduce the belongingness between the group and the remainder of the visual field. 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. But 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.

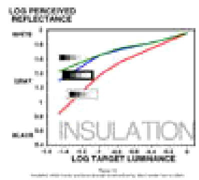

At this point this insulation effect remains little more than an empirical result; we have no deeper explanation for it. But the effect is apparently not reducible to local contrast. When the five squares with a white border were compared to a window panes arrangement in which contact between each of the darker targets and white was maximized, and a concentric arrangement in which the contact was minimized, no differences were found, as can be seen in Figure 15. the key requirement seems to be merely that the inner framework is completely enclosed by a white border.

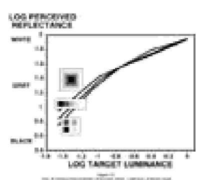

----

7. 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 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 both 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. But we intend to sidestep this issue for the moment. At the same time we will give a very concrete definition. A lightness error will be defined as 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.

Although the pattern of errors holds the potential to lead us to an adequate theory of lightness perception, the range of empirical findings on lightness errors is so vast that it is hard to know where to begin. Our approach will be to organize lightness errors into two broad classes, illumination-dependent errors and background dependent errors, corresponding to failures in the two basic kinds of constancy. A crucial test for any theory of lightness errors is to find a single model that can be applied to both of these classes of constancy failure. A canonical stimulus configuration will be taken to represent each class; the Katz experimental arrangement and the textbook version of simultaneous lightness contrast.

7.1. Gilchrist (1988): Failures due to classification problem

Gilchrist proposed a common explanation for illumination-dependent errors and background dependent errors based on edge classification. This approach is problematic because edge classification, at least for coplanar edges, is all-or-none while 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 16.

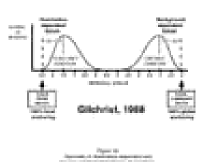

Gilchrist noted that this pattern in the data could be the result of a compromise between two competing processes, one process that produces ratio matches and one that produces luminance matches. The first of these would be invoked if the edge between the backgrounds were classified as an illuminance 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 classified to some small degree as a reflectance border. Likewise the deviation from luminance matching would also be predicted if the border between the adjacent black and white backgrounds were not altogether classified as a reflectance edge. This explanation has the appeal of symmetry. It is fully consistent with the logic of an invariant relationship between perceived surface lightness and perceived illumination level, implicit in both the albedo hypothesis (Woodworth, 1938) and the intrinsic image models described in Section 1.7.

7.1.2. Staircase Gelb data inconsistent

The classification failure model does not work, however, for the staircase Gelb effect. If illumination-dependent errors are the result of a failure to fully classify the illuminance border between two regions of different illumination it implies that the failures should be equal in magnitude for all surfaces, regardless of the reflectance of each. For example, if the illuminance step at the border of a field of special illumination is underestimated (and some portion of the step processed as a reflectance step) every surface within that field will appear erroneously lighter - but all by the same amount. However, in our staircase Gelb effect we found that each of the five squares shows a different degree of erroneous lightening. Thus the classification failure model cannot explain the data we obtained in the staircase Gelb experiments. Indeed, Bruno (1994) has reported other problems with the classification failure model as well.

7.1.3. Anchoring model consistent with Gilchrist (1988) data.

Although the staircase Gelb effect results cannot be explained by Gilchrist’s classification failure hypothesis, the results of Gilchrist’s 1988 experiments can be explained by the anchoring model. Under those conditions 100% local anchoring would produce ratio matching while 100% global anchoring would produce luminance matching. The deviations from these two poles would result from a compromise between local and global anchoring, without any reference to edge classification.

Could it be that the solution to the errors problem and the solution to the anchoring problem are one and the same? First, we must see if the local/global approach can be applied, not merely to our staircase Gelb data, but to the whole range of illumination-dependent errors.

7.2. Source of Errors

The claim implicit in this model is that errors in lightness perception are caused fundamentally by the process of compromise between local and global anchoring. In general, illumination-dependent errors are attributed to the inappropriate influence of the global (or superordinate) framework while background-dependent errors (as in simultaneous lightness contrast) are attributed to the inappropriate influence of the local framework. Any factor that strengthens the local framework relative to the global will tend to increase background-dependent errors but decrease illumination-dependent errors. Conversely, any factor that tends to strengthen the global framework at the expense of the local will increase illumination-dependent errors but decrease background-dependent errors.

8. The Model Applied to Illumination-Dependent Errors

Illumination-independent constancy is the classic form of lightness constancy. The standard experimental paradigm for studying this kind of constancy was created by Katz (1935). It consists of two side by side fields, separated by a screen so that one side is shadowed while the other side is illuminated, as illustrated in Figure 17. A target surface such as a disk is presented within each field and various psychophysical methods are used to determine which shade of gray on the shadowed side appears equal in lightness to which shade of gray on the lighted side. Our staircase Gelb effect is an illumination-independent constancy situation in the sense that the local framework represents the illuminated region and the outer framework represents the shadowed region. This differs from the Katz paradigm however, in that here the lighted region is totally surrounded by the shadowed region. In the Katz paradigm the lighted and shadowed regions are side by side. When applying our model to this paradigm we would consider the shadowed region to be one local framework, the lighted to be another, and the two together to comprise the global framework.

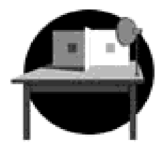

Katz, 1935

To determine how well the anchoring model can be applied to illumination-dependent errors in general, let us survey what is known about these failures in the literature. We will consider five factors that have been shown to affect the size of errors.

8.1. Errors associated with level of illumination

The fundamental form of illumination-dependent constancy failures is underconstancy, not overconstancy (Brunswick ratios under, not over 100%). Underconstancy implies either that surfaces in the brightly illuminated regions tend to appear lighter gray than they really are or that surfaces in shadowed regions tend to appear darker gray than they really are, or both. This follows from local/global anchoring. To make this concrete, consider a Katz-type constancy experiment in which a white target in the shadow has the same luminance as a black target in the lighted field. Complete global anchoring is equivalent to luminance matching, and luminance matching equals zero lightness constancy. Complete local anchoring would produce 100% lightness constancy, as long as the local framework contains a real white. Therefore any local/global compromise can yield only constancy values between zero and 100%.

8.2. Errors associated with reflectance of target.

Perhaps it has been assumed that illumination-dependent constancy failures apply equally to all shades of gray; that, for example, when surfaces in a lighted region appear lighter than their true value, all shades of gray are lightened by the same degree. However, our anchoring model predicts that illumination-dependent lightness errors are not uniform for targets of different reflectance within the same framework. The model predicts a gradient of error: in bright illumination, the lower the target reflectance, the greater the error. In shadow the reverse is predicted. This gradient of error follows from the fact that the luminance range in the global framework is greater than that of any local framework. As a result the global lightness assignments for a group of surfaces will be compressed relative to their local assignments.

There are few data that can be used to test this prediction because traditional techniques for measuring constancy allow only a measurement of the overall constancy failure, not the specific failure of a particular target. For example, in Katz's original lightness constancy experiment, the task for the observer was to adjust the gray level of a disk standing in one field of illumination so that it appears equal in lightness to another disk standing in the neighboring field of illumination. The problem is that when the match is made the experimenter doesn't know the perceived shade of gray of either disk, only that the two disks appear to be the same shade of gray. Testing for a gradient of error requires matches made with a Munsell chart.

The data from our staircase Gelb experiment (Cataliotti & Gilchrist, 1995) which were collected using a Munsell chart, show a pronounced gradient of error, as can be seen in Figure 8.

Gilchrist (1980) tested lightness constancy using different depth planes to produce the lighted and shadowed fields and the data were taken using a Munsell chart. Each plane contained two surfaces: one white and one black. In the lighted plane, the black showed an erroneous lightening while the white showed no error, while the reverse occurred in the shadowed plane. This is the same gradient of error predicted by the model.

8.3. Errors associated with reflectance of backgrounds

Helson (1943) conducted a series of lightness constancy experiments using Katz's side-by-side fields of illumination and shadow illustrated in Figure 17. White, gray, and black backgrounds were used. Constancy was best for the white backgrounds (though not great there either), worse for the gray, and poorest for the black. The same pattern of results was obtained by Kozaki (1963, 1965) as well. Leibowitz, Myers, and Chinetti (1955) tested constancy for a middle gray target on a background that was either white, gray, or black. They obtained almost zero constancy, that is luminance matching, for the gray and black backgrounds, but much better constancy for the white background. This finding of greater constancy on backgrounds of higher reflectance has remained simply a curiosity in the literature. It does not flow from any current theories, nor have any explanations been offered, as far as we know.

This finding does flow from the asymmetry inherent in anchoring by white (as opposed to black). With dark gray backgrounds, each target will usually be the highest luminance in its local framework, and will have a local assignment of white, regardless of its actual reflectance. Thus the targets will have the same local assignments even when they differ substantially in reflectance. In effect, this throws the decision to global anchoring, which finds a match when target luminances are equal; that is, zero constancy.

To be concrete, imagine that a light gray target on the shadowed side and a dark gray target on the lighted side have equal luminance values. They will have the same global assignments and, assuming each is the highest luminance in its local framework, they will have equal local assignments (white). Thus the two targets will be perceived as the same lightness, but this represents zero constancy.

8.3.1. The key: Increments versus decrements

But in fact, the special role of the highest luminance in the new model implies that the reflectance of the background is not the crucial variable, per se; but whether the targets are increments or decrements. This leads to the prediction that a pair of incremental targets on a background of higher reflectance could show less constancy than a pair of decremental targets on a background of lower reflectance. By varying both the reflectance of the background and the reflectance of the targets, Kozaki (1963, 1965) has shown that this is indeed the case.

8.4. Errors associated with articulation level

It has been shown that good constancy requires rich articulation and that pronounced failures of constancy occur when a field of illumination is poorly articulated.

8.4.1. Burzlaff experiments.

Although there has been virtually no serious discussion of this concept in the past 50 years, there is an ample amount of empirical evidence for the effectiveness of the concept of articulation. One can find several publications in the early literature that deal with the topic. A good example is the elegant work of Burzlaff (1931). Burzlaff conducted two lightness constancy experiments, one with low articulation and one with high articulation. His first experiment was modeled after one of Katz's basic techniques for studying lightness constancy. Observers were presented with two gray disks. One was mounted in a position relatively near a light source; the other was placed some distance away from the light source where the illumination level was 20 times lower (see Figure 18). The task of the observer was to adjust the gray level of one of the two disks to make it appear the same shade of gray as the other disk. Like Katz before him, Burzlaff obtained very poor lightness constancy under these conditions.

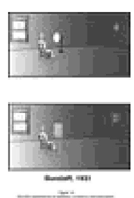

But in another experiment, Burzlaff presented a cardboard panel containing 48 chips of different shades of gray spanning the entire black/white scale, arranged in ascending order, in the position near the light source, and a second panel containing the identical 48 patches arranged in a haphazard fashion in the position distant from the light source. This is probably the earliest use of a Mondrian pattern in the literature.

The task of the observer was to select a patch from the near panel that appeared to be the same shade of gray as a given target patch in the far panel. Under these conditions Burzlaff obtained excellent lightness constancy, with hardly any measurable failure of constancy at all.

Consistent with local and global anchoring, the higher degree of articulation in Burzlaff's cardboard panel (as compared with the single disk) made the local frameworks much stronger (or, what is equivalent, increased the belongingness of each surface in the local framework to every other surface in that framework). Strong local anchoring and a real white in each local framework produces veridical perception. In the case of the single disk, local anchoring is weak and lightness is strongly determined by global anchoring. Global anchoring produces luminance matching and, in constancy experiments this means poor constancy.

8.4.2. Arend & Goldstein experiments.

Arend and Goldstein (1987; see also Arend & Spehar, 1993a and 1993b) simulated on a CRT screen either two disk/annulus patterns, one in bright illumination and one in low, or two Mondrian patterns. Observers adjusted the luminance of a target in one illumination to match, for either lightness or brightness, a target in the other illumination. The data fell closer to ratio matching for decremental targets (Section 8.3.1), for the lightness task (Section 10.6.1), and for the Mondrian stimulus (Section 6.3.2), but closer to luminance matching for incremental targets, for the brightness task, and for the disk/annulus stimulus. All of these are consistent with the anchoring model. In these experiments, luminance matching is equivalent to global anchoring while ratio matching is equivalent to local anchoring. The higher articulation level of the 32-patch Mondrian, compared to the 2-patch disk/annulus pattern, produced almost completely local anchoring.

8.4.3. Schirillo experiments.

Schirillo, Reeves, and Arend, (1990) published a replication of one of Gilchrist's (1980) experiments on lightness and depth perception. Using a CRT screen as a stereoscope they replicated the basic Gilchrist finding but with larger failures of constancy than Gilchrist obtained. In a follow-up paper, Schirillo and Arend (1995) investigated the cause of these failures of lightness constancy. They hypothesized that these errors are due to local contrast effects caused by non-coplanar surfaces that are nonetheless retinally adjacent to the target surface. They showed that the failures of constancy are indeed greater when the target surface is situated on the border between the differently illuminated planes than when the same surface is embedded within the plane. We would attribute this to the factor of retinal proximity mentioned in Section 5.5.

But even for their target on the illumination border, subjected to a local contrast effect, the degree of constancy was far higher in the second paper than in the first. The reason for this seems clearly to do with articulation, as shown in Figure 19. In the first experiment there were only 2 or 3 surfaces in each plane, while in the second experiment one plane contained a Mondrian composed of 46 surfaces, the other plane a Mondrian of 15 surfaces.

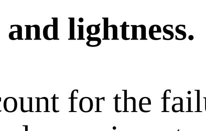

8.4.4. Application of model to depth and lightness.

In summary, our proposed model would account for the failure of constancy (and the failure of the coplanar ratio principle) in Gilchrist's original experiment and in the Schirillo, Arend, and Reeves experiment in the following way. Because each plane has low articulation, each constitutes a relatively weak framework and the global relationships intrude substantially to produce failures of constancy. The white surface in the dimly illuminated near plane, for example, would have a global value somewhere around middle gray and this would cause it to appear somewhat darker than a good white, which it did. Likewise, when the same target appears to lie in the far plane, it would be white in the local framework but again roughly middle gray in the global, causing it to appear somewhat lighter than black, which it did. When the local frameworks are articulated they become much stronger, resisting global anchoring, thus producing greater constancy.

Work by Wishart, Frisby, and Buckley (1995) shows that the strength of the coplanar ratio effect does vary with its degree of articulation. But because they measured brightness, and not lightness, this work is discussed later together with a class of stimuli that produce background-dependent errors.

Other work showing an effect of articulation has been reported by Cramer (1923), Henneman (1935), and Katona (1929).

8.5. Errors associated with field size

MacLeod (1932), in a constancy experiment modeled after an ingenious Kardos experiment, tested the perceived lightness of a gray disk located in the center of a cast shadow of variable size projected onto a white wall. As the shadow grew larger the disk appeared lighter gray and thus constancy increased, just as the law of field size claims. According to our anchoring model, the penumbra at the border of the shadow isolates the shadowed region creating a local framework consisting of two regions, the gray disk and its white surround. Anchored within this framework the disk appears gray. But anchored more globally, to the brighter white wall outside the shadow the disk would appear much darker. As the shadow gets larger, the local framework that it represents gets stronger and the gray disk becomes anchored more strongly to its local surround, causing it to lighten.

Agostini & Bruno (1996), Henneman (1935), and Katona (1929) have also reported effects of field size consistent with our anchoring model.

9. The Model Applied to Background-Dependent Errors (Simultaneous Lightness Contrast)

There is an obvious structural similarity between the simultaneous lightness contrast display and the Katz paradigm with side-by-side lighted and shadowed regions. Each is organized into two side-by-side subordinate frameworks, which together compose a single superordinate framework. For continuity in the use of the terms local and global we begin by applying our model to the case in which the contrast display fills the observer's entire visual field. But we must keep in mind that more typically (as in a textbook, for example) the contrast display occupies only a part of the visual field. In that case we must use the more cumbersome terms subordinate framework and superordinate framework. The global framework in that case would include both the contrast display and its surrounding context, changes in which can have a pronounced effect on the contrast illusion, as Agostini and Bruno (1996) have shown.

According to our anchoring model, there is both an anchoring component and a weaker scaling component in the contrast effect. We begin with the anchoring component that we believe accounts for the larger part of the illusion.

9.1. Anchoring component

In the global framework, the two targets have the same assigned lightness value because they have the same luminance and thus they share the same luminance ratio with the global anchor (refer to Figure 20), which in this case is the white background on the right side. But in the two local frameworks the situation is as follows. Because the local framework on the right includes the global anchor (that is, the white background) the target on the right stands in the same luminance ratio to the anchor, both locally and globally. Anchoring per se doesn’t alter this target. The target on the black background, on the other hand, has a local assignment that is very different from its global assignment. Locally the target is the highest luminance so that its assignment is white. Thus a compromise between global and local assignments yields a higher value for the target on the black background than for the target on the white background, how much higher depending on the local/global weighting.

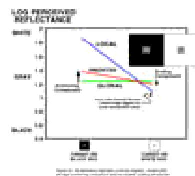

Virtually the same analysis of the simultaneous contrast display was offered in 1987 by McCann: "If global normalization of the entire field of view were complete, we would expect that observers would report the two gray squares with identical reflectances would have the identical appearance. If local mechanisms were the only consideration, then the gray square in the black surround should mimic the results found in the Gelb experiment, and should appear a white, since it is the maximum intensity in the local area. Observer results give important information about the relative importance of global and local interactions. The gray square in the black surround is one lightness unit out of nine lighter than the same gray in the white surround. If local spatial calculations were the only consideration, the gray in black should appear a 9.0. If global spatial considerations were the only consideration, the gray in black should appear a 5.0. The observer matched the gray in black to a 6.0. In other words, the spatial normalization mechanism is an imperfect global mechanism. Alternatively, it is a local mechanism that is significantly influenced by information from the entire image." (On the Munsell scale, 9.0 is white and 5.0 is middle gray.)

9.2. The scale normalization component

Because each of the local frameworks in the contrast display contains a limited luminance range, we must expect a small scale normalization effect, or expansion, that, because the anchor is at the top of the lightness scale, is expressed as a darkening of the darker region in each framework. This means that the only target affected is the one on the white background.

9.3. Predictions of the model

The above analysis leads to at least the following four testable predictions of the model:

9.3.1. The locus of the error.

In our work on simple or single-framework images we found that the anchoring effects are much stronger than the scaling effects. Thus if our assumption is valid that local frameworks embedded in complex images function qualitatively like single-framework images, we must also expect that the anchoring component of our model is stronger than the scaling component. Combining this with the fact that the anchoring component produces a lightening of the target on the black background (but not a darkening of the other target) while the scaling component produces a darkening of the target on the white background (but no lightening of the other target), we derive the prediction that most of the error represented by the simultaneous contrast effect occurs for the target on the black background.

It is difficult to determine how much the simultaneous contrast illusion involves a lightness distortion of the target on the black background and how much the target on the white background. The difficulty stems from the lack of a completely valid measure of perceived lightness. Matches made from a standard Munsell chart indicate that the distortion occurs only for the target on the black background, as the new model predicts. However, the chips in the standard Munsell chart are presented on a white background. Presumably a Munsell chart with chips on a black background would suggest that the distortion occurs for the target on the white background.

Gilchrist and Economou (unpublished) obtained lightness matches for the simultaneous contrast display using a Munsell chart on a black and white checkerboard background so that each chip borders black and white equally. The illumination on the Munsell chart was equal to that on the contrast display. They conducted three variations of the experiment, using a different group of ten subjects in each condition: (1) both contrast display and Munsell chart on paper, (2) both on CRT, and (3) contrast display on CRT, Munsell chart on paper. The results are shown in Figure 21. Consistent with the anchoring model, the target on the black background consistently showed the main error, with a smaller error for the other target.

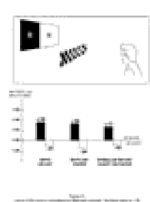

9.3.2. The size of the error.

Our model predicts that the size of the contrast illusion depends on the relative strengths of the local and global frameworks. The local frameworks are relatively weak because they have minimal articulation. As for the global framework, we have so far applied the model to the case in which the contrast display fills the entire visual field, even though there are no published reports of the illusion under these conditions. More typically the contrast display is presented within a larger complex image. In that case the strength of the framework that consists of the contrast display itself depends on how strongly that framework belongs to the larger image and how strongly it is segregated from the larger image. When the display is presented in a textbook, it is perceived to belong to the page of the book and to the table on which the book is lying. Thus the global, or superordinate, framework is quite strong, so the illusion should be quite weak. This is just the typical result: the target on the black background appears only slightly lighter (by about 3/4 of a Munsell step) than the target on the white background.

However, we would expect the illusion to be strengthened if the entire contrast display is perceptually segregated from the global framework. This outcome has been shown to occur by Agostini and Bruno (1996). They showed that a substantially stronger illusion occurs if the contrast display is illuminated by a spotlight so that the bright illumination perceptually segregates it from the less-brightly illuminated surrounding region. Under these conditions the illusion is twice as great, 1.5 to 2 Munsell steps, as when presented under more typical conditions, such as in a textbook lying on a table within a lighted room. As the scope of the spotlight is gradually enlarged to include more of the surrounding region the strength of the illusion is gradually reduced to 3/4 of a Munsell step. The same result occurs when a shadow is cast that falls only on the contrast display.

Here we see another example of Katz’s law of field size. As the size of the spotlight increases, greater constancy results.

Agostini and Bruno also found that presenting the SLC display on a CRT screen has the same effect as presenting it within a spotlight; the illusion is approximately doubled. This is an important finding, given the need to integrate the older lightness literature based on paper displays with the more modern work done primarily on CRT screens. The implication of their finding is that even the perceived lightness in a region of the visual field the size of a CRT screen depends ultimately on coexisting stimuli in the larger context.

The model predicts that the contrast illusion can be strengthened by increasing the articulation of the local frameworks. A high articulation version of simultaneous contrast is presented in Figure 22c. Here the white background has been replaced with light gray patches and the black background with dark gray patches. The illusion is strengthened even though the two backgrounds now differ less in space-averaged luminance. Adelson (1996) has reported the same effect.

9.3.3. No contrast effect with double increments

We noted earlier that, according to our model, the source of the contrast effect lies in the local assignments, not the global because they have equal luminance, and thus equal global assignments. This implies that if local assignments were also somehow equal the illusion would be eliminated. The simultaneous contrast display with double increments provides just such a case and, as can be seen in Figure 22, it seems to show almost no contrast effect at all (as reported by Gilchrist, 1988) when compared with the standard display. We can further predict that no illusion will occur when the double increments version is presented within a spotlight, except that now both targets should appear white. These findings have recently been confirmed in our lab.

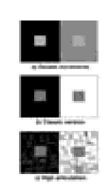

It may be noted that the above application of the model to the simultaneous contrast pattern is essentially identical to the application of the model to illumination-independent constancy using black backgrounds outlined in Section 8.3. Thus the model predicts that the same double increment conditions that produce poor constancy with respect to different illumination levels, produce good constancy with respect to different backgrounds.

9.3.4. Segregating effect of luminance ramps

Various manipulations of the contrast display support the idea that luminance ramps exert a segregative role, in belongingness terms. It has been generally known that if the sharp black/white boundary of the contrast display is replaced by a gradient, the illusion is strengthened. By segregating the two halves of the display we can say either that the superordinate framework is weakened, or its equivalent, that the two local frameworks are strengthened. Agostini and Galmonte (1997) have recently shown a particularly strong version of this illusion.

9.3.5. Belongingness creates the illusion

According to our model, any of the grouping principles should suffice to create a contrast effect. This can be seen in a variety of published reports.

9.4. The Benary effect

The Benary effect (1924), illustrated in Figure 23, was offered by Wertheimer as evidence of the crucial role of belongingness. The gray triangle that appears to lie on top of the black cross appears lighter than the gray triangle that appears to lie on the white background, despite the fact that the two triangles have equivalent surrounds. Each triangle borders white on its hypotenuse and black on the perpendicular sides. In terms of belongingness, however, when perceptual structure is taken into account, one triangle appears to belong to the black background and it appears lighter than the other triangle that appears to belong to the white background.

9.5. White's illusion.

Strong support for the role of belongingness in contrast effects comes from White's illusion (White, 1981), shown in Figure 24. Here two targets of equal luminance appear different but in a direction opposite to what a contrast model would predict. The gray bars that appear lighter are those that are surrounded mostly by white while the ones that appear darker are mostly surrounded by black. Yet one clearly has the phenomenal impression that the lighter gray bars belong to the black stripes, not to the white stripes. The main grouping factor at work here is the T-junction.

9.6. Checkerboard contrast.

The challenge to a lateral inhibition account is pushed to the limit in the phenomenon of checkerboard contrast, shown in Figure 25 as a variation on a figure presented by DeValois and DeValois (1988, p. 229). Here the gray square that appears lighter is completely surrounded by white; the square that appears darker is completely surrounded by black. In our view this illusion results from the salient grouping along the diagonals. The square that appears lighter may be completely surrounded by white, but it is perceptually grouped with, and thus anchored by, the black squares.

The effect is rather weak to be sure, and somewhat unstable. But we argue that this is just what should be expected when a fuller accounting of the grouping factors is made. Indeed this stimulus provides a good illustration of multiple local frameworks.

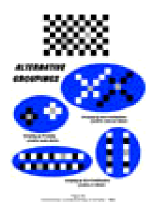

Each target is a member of three kinds of groups, but each kind of group produces a different effect on lightness. Grouping by rows and columns (based on good continuation) produces no differential lightness effect because the highest luminance within groups is the same for both targets. Grouping along the diagonals, based also on good continuation, produces the main effect. Finally our model always postulates a weak effect of grouping by retinal proximity. This effect, which is the conventional simultaneous contrast effect, opposes the effect produced by diagonal grouping. The weak effect we observe equals the effect of grouping by diagonals minus the effect of grouping by retinal proximity.

9.7. Agostini and Proffitt: Common fate.

Agostini and Proffitt (1993) demonstrated the importance of belongingness for simultaneous contrast by eliciting the effect in a novel way. A flock of black disks, a flock of white disks, and two gray disks were placed randomly on a blue background. Not surprisingly the two gray disks appeared the same shade of gray. However, when all of the black disks and one of the gray disks were caused to move in one direction, and all of the white disks and one of the gray disks were caused to move in a different direction, as shown in Figure 26, a weak simultaneous contrast effect was produced. The gray disk moving with the black disks appeared lighter than the gray disk moving with the white disks.

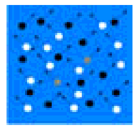

COMMON FATE
Agostini & Proffitt, 1993

This result shows that common fate causes one gray disk to belong to the group of black disks and the other to belong to the group of white disks. Thus two overlapping frameworks are created and the gray disk that belongs to the framework composed of black disks appears slightly lighter because it is the highest luminance within that framework. The lateral inhibition account of simultaneous contrast fails completely here because the two gray target disks have the same blue background.

9.8. Laurinen and Olzak.

Laurinen and Olzak (1994) superimposed shallow sinusoidal modulations on each of the four parts of the simultaneous contrast display. They found that the illusion is substantially weakened if in both local frameworks the modulation frequency on the target is different from that of its background. Our interpretation of these findings is that the sinusoidal modulation provides the basis for grouping by similarity. The difference between each target and its background segregates the target from its background, weakening the local framework, which we regard as responsible for the illusion. At the same time the similarity of left and right targets increases their belongingness as does the similarity of left and right backgrounds, with the effect of strengthening the global (or more correctly, the superordinate) framework.

9.9. Wolff illusion

Wolff (1934) created the patterns shown in Figure 27 to illustrate the role of figure/ground segregation in lightness. Each large disk contains equal retinal areas of light and dark. The light gray appears lighter when it appears a figure than it does as ground, and the dark gray appears darker as figure than as ground. But the background, because it appears to continue behind the figures, has a larger perceived area than the figures. Accordingly the pattern on the left comes under the Area Rule and consistent with the schematic in Figure 6, both light and dark regions should become lighter. This is just what happens.

We can see the local/global compromise at work here as well. The illusion would be much greater if each of Wolff's patterns were painted onto the inside of a dome so that it filled the whole visual field, as Figure 5 confirms. Placing the two patterns on a page, as in Figure 27 creates a superordinate framework within which there would be no illusion. What we see in Figure 28 is a compromise between no illusion and the stronger illusion that would occur in the single framework of a dome.

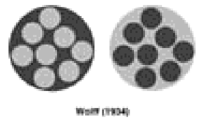

9.10. Depth manipulations:

Given the importance of coplanarity as a factor in belongingness, one should be able to reduce the size of the simultaneous contrast illusion by moving the target squares to a depth plane midway between the observer and the black and white backgrounds. As long as the black and white regions continue to form the retinal backgrounds for the two targets, the illusion should be unchanged according to the conventional interpretation based on lateral inhibition. But according to our model this depth manipulation should decrease the belongingness between each target and its background, weakening the two local frameworks, which should in turn weaken the contrast illusion. Wolff (1933) conducted such experiments, concluding that (p. 97): "contrast is strongest when both fields lie in the same plane and there is no contrast at all when the fields are phenomenally situated at a large distance from each other." However, it should be noted that Gibbs and Lawson (1974) found no such effect.

Several investigators have reported that White’s illusion can be reversed using stereopsis to alter the belongingness relations (Anderson, 1997; Spehar, Gilchrist & Arend, 1995; Taya, Ehrenstein & Cavonius, 1995)

9.11. Adelson’s corrugated Mondrian

Adelson has created a delightful illusion, illustrated in Figure 28, based on the perceived planarity of surfaces. Surface B appears lighter than surface A even though they have the same luminance value. According to Adelson (1993, p. 2044), surface A is seen as "a dark gray patch that is brightly lit" while surface B is seen as "a light gray patch that is dimly lit." This interpretation of the illusion is consistent with that of Gilchrist (1980, Gilchrist, Delman, & Jacobsen, 1983), although both Gilchrist and Adelson have moved away from this construction.

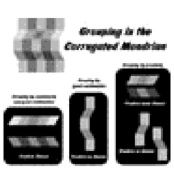

This illusion, which is obviously quite strong, follows directly from our anchoring analysis. And again, as in checkerboard contrast, the source of the strength seems to lie in the multiple grouping factors. There are three obvious groupings of the targets: (1) horizontally by rows, (2) vertically by columns, and (3) by local retinal adjacency. Grouping by columns produces no effect (targets share the same group and the same highest luminance). Grouping by rows produces the main effect. Target B has a higher local lightness assignment than target A because the highest luminance in the row is lower for target B than for target A. Some of the possible groupings by local retinal adjacency produce a weak effect in the same direction as grouping by rows. Thus for the corrugated Mondrian, grouping by local adjacency augments the main effect of grouping by rows whereas for checkerboard contrast it subtracts from the main effect of grouping by diagonals.

Although Adelson's results are stronger than simultaneous contrast and checkerboard contrast, they are much weaker than those of Gilchrist (1980). This is not surprising from the viewpoint of our model. Adelson presented figures on a CRT screen while Gilchrist used paper 3D displays; both used binocular viewing. Thus in addition to pictorial grouping factors we must also consider grouping factors inherent in stereopsis. In the Gilchrist experiments, stereo factors can be expected to further segregate regions by planarity while in the Adelson experiments, these same stereo factors work against planarity segregation produced solely by pictorial cues.

9.11.1. Wishart experiments.

Wishart, Frisby, and Buckley (1995) have shown that the size of the illusion varies as a function of the perceived angle of the folds, as shown on the left side of Figure 29. Gilchrist (1977, 1980) treated coplanarity in an all-or-none fashion, but Wishart et al show that the degree of belongingness between perceived planes varies with the angle between them.

If we can assume that planarity similarity is a graded variable, then all the main features of Wishart's data follow from the new model. (1) As the angle of the fold becomes sharper, the belongingness between the perceived planes weakens, weakening the strength of the superordinate framework (the whole display). The data shift closer to the local assignments (see Figure 29) increasing the strength of the illusion. (2) The change is expressed by target B, not A, because the highest luminance in target A's main local framework (its row) is the same as the highest luminance in the superordinate framework, thus a change in local/global weighting doesn't change the predicted lightness of A.

(3) In the flat condition the illusion is weakened because planarity similarity both strengthens the vertical column and weakens the horizontal row. (4) In the vertical condition the illusion is weakened even further because the column is further strengthened and the rows are further weakened. The residual illusion in the vertical condition can be attributed to the weak effect of retinal adjacency.

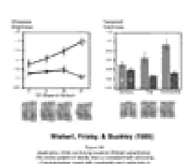

These same investigators found that the illusion is absent in the simpler display shown in Figure 30. If the basic Adelson effect is based on the perceived difference of illumination in the planes of surface A and surface B, the illusion should still be present. However, according to our model, the strength of the local framework for surface B should drop to almost zero because its degree of articulation has the minimal value of only one. Thus we would expect the percept to be strongly dominated by the global assignments (A and B equal) and this is what is obtained.

----

10. Brightness Induction: Contrast or Anchoring?

During the 1950s and 1960s an extensive literature developed under the rubric of brightness contrast or brightness induction. The goal of this industry was to explain the perception of brightness and, by extension, surface lightness, with the physiological mechanism of lateral inhibition. These experiments were very similar to each other in many respects. The stimuli consisted of homogeneous bright patches presented within a completely dark context. The test display usually consisted of a target square or disk, known as the test field, and an adjacent, nearby, or surrounding region known as the inducing field. The inducing field was usually brighter than the test field. In addition there was a separate region located elsewhere in the visual field, often presented to a separate eye in haploscopic presentation, that served as a matching region and was adjusted by the observer to match the perceived brightness of the test field. Three independent variables were tested repeatedly (1) the luminance, (2) the size, and (3) the proximity of the inducing field. The results of these experiments were interpreted within the theoretical framework of contrast (Hering, 1874). But we believe the results make more sense as manifestations of anchoring. The first two variables, luminance and size, can be accounted for by rules of anchoring in simple images. The third, proximity, involves the concept of a local/global competition.

10.1. Induction

Perhaps the best known of these experiments is that of Heinemann (1955), who varied the luminance of an annulus surrounding a disk of fixed luminance. For every setting of the annulus the appearance of the test field was measured by having the observer adjust a separate, isolated disk so as to match the appearance of the test disk. The basic pattern of data he obtained are shown on the left side of Figure 31.

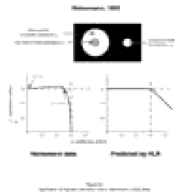

There are several features that characterize this curve. First there are two main sections, one roughly horizontal and one roughly vertical. The break point in the curve occurs roughly at the increment/decrement threshold. Heinemann found that as long as the luminance of the annulus is lower than that of the disk, annulus luminance has little effect on the appearance of the disk. However, as soon as the annulus luminance surpasses that of the disk, the disk begins to darken.

This basic shape is a manifestation of the highest luminance rule. Figure 31b shows the predicted shape of the curve under the highest luminance rule. Until annulus luminance supersedes disk luminance, changes in annulus luminance have little effect on the perceived lightness of the disk, which always being the highest luminance continues to appear white. Perception of the test disk begins to change only when the annulus becomes highest luminance in the display. From then on, further increases in annulus luminance produce corresponding (that is, consistent with Wallach's ratio principle) decreases in disk lightness.

But when this theoretical curve is superimposed on Heinemann's empirical curve three discrepancies are revealed. First, anchoring predicts a drop in the curve with slope -1. Heinemann's curve drops much faster because the comparison disk is presented within a totally dark background. This means that the comparison disk can never appear as a decrement; it can never be set to a luminance value below that of the surround. Whittle (1994a, 1994b) has noted that matching an increment to a decrement is extremely difficult and subjects will not make such a match if they have the chance to match a decrement with a decrement or an increment with an increment. Further, as Katz (1935) made very clear, although a disk surrounded by an annulus can appear in the surface color mode, a disk surrounded by darkness can appear only in the aperture color mode (as a luminous disk). Matches across these modes are not reliable. As soon as the test field in Heinemann's experiment becomes a decrement it begins to appear gray and the curve drops precipitously as the observer makes a futile attempt to escape the luminous appearance.

When matching is done with a disk/annulus pattern or a Munsell chart, the slope of -1 is obtained. Wallach (1948), Heinemann (1955) himself, and Gilchrist and Bonato (1995) obtained such a function.

2. Enhancement effect. The first section of the curve is not totally horizontal, as the highest luminance rule demands, but has an upward bulge that has been referred to as the enhancement effect. MacLeod (1947) reported this effect as well.

3. Breakpoint offset. The breakpoint (where the curve drops below its initial y-value) does not always occur exactly at the increment/decrement threshold (the point where the annulus becomes the highest luminance), as the highest luminance rule predicts. In Heinemann's study, it occurs a little before that point. In other studies (Heinemann, 1971) it occurs even earlier while in still other studies (Diamond, 1953; Horeman, 1963; Leibowitz, Mote, & Thurlow, 1953; Torii & Uemura, 1965) it occurs right at the increment/decrement threshold.

Two important facts that have already been established provide the key to understanding these features of the data. First, the breakpoint offset and the enhancement effect appear and disappear together (Heinemann, 1971; Torii & Uemura, 1965), and second, they both occur only when the Area Rule is applicable. Heinemann (1971, p. 152) notes: "...as the size of the inducing field is decreased the enhancement effect...becomes smaller." Heinemann and Chase (1995) have recently offered a mathematical model that includes an account of the enhancement effect, but neither this model nor the earlier attempt by Diamond (1960) appears to account for why the enhancement effect comes and goes with relative area.

Both the enhancement effect and the breakpoint offset, as well as their linkage, are readily understandable in terms of the Area Rule. Remember that when the Area Rule is engaged, it produces a lightening of the darker (and larger) region, even without a change in the luminance ratio between the two regions. This lightening effect has two consequences. First, it puts an "upward pressure" (see Figure 6) on the appearance of the brighter region (the disk) causing it to appear as a kind of super white, or the enhancement effect that Heinemann obtained. Second, it causes the annulus to become white even before it becomes the highest luminance. As soon as the annulus becomes the anchor, any further increases in its luminance must make other members of the framework appear to darken. Thus the curve representing disk appearance begins its dive even while it is yet an increment, which shows up as the breakpoint offset.

The enhancement effect and the breakpoint offset have been obtained only in those experiments (Heinemann, 1955; Torii and Uemura, 1965) with inducing fields of larger area than the test field. When the test field is as large as the inducing field (Diamond, 1953; Horeman, 1963; Leibowitz, Mote, & Thurlow, 1953; Torii & Uemura, 1965), neither enhancement nor breakpoint offset occurs.

10.2. Area effects

Besides accounting for Heinemann's data, the Area Rule is consistent with three other brightness induction studies.

Stevens (1967) used a disk/annulus display in a study of area and brightness that was similar to, but more systematic than Wallach's study. Diamond (1955) carried out an analogous study except that his stimulus consisted of a pair of adjacent rectangular regions. Both Stevens and Diamond found a pronounced effect of relative area on perceived lightness only when the requirements of the Area Rule were satisfied; that is, only when the region of lower luminance had the higher relative area. The data presented by Stevens and Diamond are reproduced in Figure 32. Shading has been added to both graphs to indicate the zone in which the Area Rule applies, showing that effects of relative area on lightness/brightness occur primarily within this zone. Outside of this zone, relative area has little or no effect on perceived lightness.

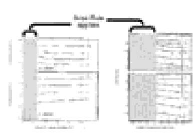

In another variation of the Gelb effect, Stewart (1959) examined how the area or size of a relatively small but relatively bright disk affects the perceived brightness of a larger but darker disk upon which it is placed. All of the conditions that Stewart tested fall within the applicable zone of the Area Rule, and he obtained a pronounced effect of area on brightness, consistent in direction with the Area Rule.

10.3. Separation

The effect of separation between inducing field and test field on the appearance of the darker test field (holding the luminance of each constant) has been studied by Cole and Diamond (1971), Dunn and Leibowitz (1961), Fry and Alpern (1953), and Leibowitz, Mote and Thurlow (1953). All found that the perceived brightness of the darker test field increases with increasing separation between the two. Again, no change occurs in the brighter region. Both McCann and Savoy (1991) and Newson (1958) obtained analogous results using similar stimuli and measuring lightness, not brightness. Most of these investigators (Newson and McCann & Savoy are the exceptions) gave a lateral inhibition account of their results, according to which the test field becomes brighter with increasing separation due to the well-known drop off of inhibition strength with spatial separation across the retina.

But we argue that these results are better explained in terms of anchoring. The brighter of the two doesn't change in brightness because, whether it is anchored relative to the dark surround or anchored relative to the darker target, it has the highest luminance and thus it appears as a constant white. As for the darker region, the greater its proximity to the brighter region, the more it belongs to the brighter region and the more its lightness value is determined by its relationship to the brighter region. The less its proximity to the brighter region, the more its lightness is determined relative to the dark environment, in relation to which its assigned value is white.

There are two lines of research that have attempted to test between contrast and anchoring interpretations of the separation effect, one using depth separation and one using coplanarity.

10.4. Depth separation: Gogel and Mershon

Rock's (1975, 1973) classic question, that we applied earlier to the area factor, can be applied to the separation factor as well. Is brightness determined by retinal proximity or by perceived proximity? Lateral inhibition implies that retinal proximity is crucial while the concept of proximity in three-space is more relevant for the perceptual grouping used in our anchoring model. Gogel and Mershon (1969; Lie, 1969; Mershon, 1972; Mershon & Gogel, 1970) showed that when two adjacent surfaces presented within a dark room are separated in stereoscopic depth while maintaining their retinal adjacency, the perceived lightness of the darker region moves toward white just as it does when the two regions are separated in the frontal plane. This is consistent with Rock's earlier finding (Rock and Brosgole, 1964) that grouping by proximity depends on perceived proximity, not retinal proximity. And it supports an anchoring account over a lateral inhibition account.

10.5. Gelb effect: Contrast versus anchoring

Many years ago Adhémar Gelb (1929) demonstrated that if a piece of black paper is suspended in midair and illuminated by a beam of light that illuminates only the black paper it will appear white. In this case the black paper is actually the highest luminance in the scene. But when a piece of real white paper is brought into the beam and placed next to the original black paper, the black paper becomes perceptually gray or black. Indeed the black paper is perceptually darkened by any other adjacent paper that has a luminance that exceeds that of the black paper.

Gelb's marvelous experiment pre-dated the brightness induction literature by several decades. But the similarity between Gelb's stimulus conditions and those of the induction experiments (a few brightly lit regions presented in a dark environment) guaranteed that Gelb's illusion would soon be accounted for in terms of contrast based on lateral inhibition. Explicitly contrast accounts of the Gelb effect can be found in Hurvich and Jameson (1966, p. 107) and Cornsweet (1970, p. 368). Stewart (1959) showed that the amount of darkening of the target depends on both the size and the location of the white inducing region arguing that this supports the contrast account.

Gelb did not offer an explicit anchoring interpretation of his illusion, but he did clearly reject the contrast interpretation. Koffka (1935, p. 246) also rejected a contrast interpretation, offering an analysis that proceeds within the framework of anchoring, although at times he seems to give the lowest luminance equal status with the highest luminance. McCann (1987, p. 279) has been the clearest in attributing the Gelb effect unambiguously to anchoring in general and to the highest luminance rule in particular: "Here a black piece of paper is the only object in the field of view, by itself it looks a dim white, but unquestionably white. When a white paper is put beside the black paper then white looks white and black looks black. This is an example of a global normalization process..."

Although both contrast and anchoring interpretations have appeared in the literature, there was no explicit test between them until recently. Cataliotti and Gilchrist (1995; see also Gilchrist & Cataliotti, 1994) exploited another difference between contrast and anchoring treatments of the separation factor. From a contrast perspective, the strength of proximity effects should not depend on whether or not the space between the test and inducing fields is filled with a coplanar surface. But from an anchoring perspective this can make a crucial difference. For example, moving the position of the highest luminance around within a Mondrian has little effect on lightness values because the various regions of the Mondrian are already strongly grouped together by coplanarity and adjacency.

Cataliotti and Gilchrist tested the proximity factor by presenting the Gelb effect in a series of incremental steps. Their observers were initially presented with a square piece of black paper in the beam of light. It was duly perceived as white. Then a piece of dark gray paper was introduced into the beam of light in a position immediately adjacent to the original black paper. The dark gray square was then perceived as white because it was now the highest luminance in the display, and the original black square came to appear light gray, relative to the new white. Then a middle gray square was introduced into the beam of light next to the dark gray square. At this point the middle gray square appeared white and both the dark gray square and the black square became darker perceptually. The procedure was repeated with a light gray square and finally with a true white square. At this point the stimulus consisted of a horizontal row of five squares.

The goal was to determine whether the darkening of the existing square or squares by the addition of each new and brighter square is greater in magnitude when the new brighter square is adjacent, than when it is in a more remote location. This experiment also served to test a claim made by Reid and Shapley (1988; see also Shapley & Reid, 1985) that the effectiveness of edge integration weakens with distance, much like that of lateral inhibition.

Cataliotti and Gilchrist found that the distance between the highest luminance and each target square did not have a significant effect on the darkening of the target square. This can be seen in Figure 33 in which the perceived lightness value of the original black square is plotted as a function of the maximum luminance in the group. The fact that the obtained data points fall along a straight line indicates that the perceived lightness of the black square depends solely on the luminance ratio between the black square and the square with the highest luminance, regardless of the spatial distance between them. If the darkening power of the highest luminance depended on its distance from the black square, the data points would fall along a negative slope with a positively accelerated curve.

In another experiment using the entire group of five squares, Cataliotti and Gilchrist found that it makes no statistically significant difference to the perceived lightness of the black square whether the white is placed immediately adjacent to the black square or placed at the opposite end of the row as in the staircase order.

These findings imply that the darkening effect of introducing a higher luminance into the group should be considered a matter of anchoring, not contrast.

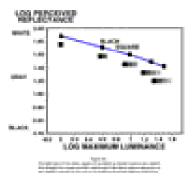

10.6. Summary: application of model to reduction conditions.

To summarize, it appears that all of the so-called contrast experiments of the 1950s and 1960s can be more consistently interpreted as evidence of anchoring rather than contrast. When two homogeneous surfaces of different luminance values are presented within an otherwise dark void, the degree to which they belong together depends strictly on their proximity. When the two are adjacent, they belong together strongly, meaning that the lightness of the darker surface is strongly determined by its relationship to the lighter surface. The greater the separation between them, the less the darker surface belongs to the lighter and the more each belongs to the surrounding darkness. Relative to this darkness, each appears white (and also luminous).

Any additional factor that also increases the degree of belongingness between the two regions will darken the appearance of the darker region. This can be achieved by: (1) totally surrounding the darker by the lighter, (2) increasing the perceived size of the two surfaces together, or (3) adding additional coplanar and contiguous surfaces. Any factor that reduces the degree of belongingness between the two regions will lighten the appearance of the darker region. This can be achieved by: (1) reducing their proximity within the frontal plane, (2) reducing their proximity in depth, or (3) reducing the coplanarity of the two surfaces, that is, changing the angle between their planes while preserving their contiguity.

10.6.1. Lightness versus brightness

The fact that the induction literature can be understood in terms of anchoring is all the more remarkable given that those experiments studied brightness (to the extent that the distinction was clear at all), and our anchoring model is a model of lightness perception. We see this as evidence that the visual system attempts to see even such reduced stimuli as surfaces rather than floating light sources. This is not to question the lightness/brightness distinction. Rather, we suggest that, under such ambiguous conditions, even when observers attempt to match targets for luminance value, perceived surface lightness intrudes, a bit like the way color names interfere with reporting the color of the ink in the Stroop test. In general, the difference between lightness and brightness instructions can be accounted for simply by changing the relative weighting of the local and global frames. Brightness instructions shift the weight to the global framework while lightness instructions shift the weight to the local framework.

11. Tests of scaling normalization effect

A truncated luminance range was used in both the Mondrian world experiment (Cataliotti & Gilchrist 1995) and the split-field dome (Li & Gilchrist, in press), and both produced expansion. This can be seen in Figure 4 by comparing the length of the "perceived" line with the length of the "actual" line.

The scale normalization effect is not large, but it does seem to account for some puzzles. For example, it explains a small part of the lightness error in simultaneous contrast, in addition to the larger anchoring effect, as noted in Section 9.2. And it explains why the luminance of a homogeneous surround that is functionally equivalent to a checkerboard surround is not exactly equal to the highest luminance in the checkerboard, but slightly lower than that (Bruno, 1992; Schirillo and Shevell, 1996). And the phenomenon of contrast-contrast (Chubb, Sperling, & Solomon, 1989) concerning the relationship between physical contrast and perceived contrast appears to be reducible to the scale normalization effect.

In some cases, however, we have found that what appears to be a scaling problem is really an anchoring problem in disguise. The compression we obtained in the staircase Gelb effect presented us initially with a serious scaling problem. But according to our anchoring model, this scaling problem reduces to an anchoring problem. All five of the squares are assigned a value of white, or close to white, in the global framework. Averaging these global assignments into the final compromise necessarily produces compression. Substantial compression effects are also found in a single framework when the Area Rule applies, as indicated in Figure 6. Here again the scaling effect reduces to a matter of anchoring. The compression is the product of a tension between anchoring by the highest luminance and anchoring by the largest area.

11.1. Crispening effect

So far we have said nothing about the crispening effect (Semmelroth, 1970; Takasaki, 1966; Whittle, 1992), perhaps because we see no obvious way to relate it to our model. But on its face the crispening effect is a case of re-scaling; there is expansion near the luminance of the background. Thus it would be reasonable to expect some systematic connection between the crispening effect and our model.

12. Problems for the Intrinsic Image Models

At the beginning of this paper we briefly traced the development of the kind of intrinsic image models that have developed during the past two decades. These models are far more effective than the earlier simplistic and wooden models based on lateral inhibition. They are much more successfully applied to a wide range of viewing conditions, especially the more complex images typical of our everyday perceptions. They more adequately capture the multi-dimensional character of visual experience reflected in our perception of both the lightness value of surfaces as well as patterns of illumination on those surfaces. Intrinsic image models (Adelson and Pentland, 1990; Barrow and Tenenbaum, 1978; Bergström, 1977; Gilchrist, 1979) are well-suited to account for veridicality in surface lightness perception.

Nevertheless, the intrinsic image models of lightness perception seem to be seriously undermined by a series of empirical challenges.

12.1. Problem of the staircase Gelb effect data.

Consider how the staircase Gelb display would be treated, according to Gilchrist's (1979) edge classification/edge integration model. In short, the large luminance range within the visual field would signal the need to parse the image into different fields of illumination. This requires identification of the illuminance edges. The basis of edge classification may be ambiguous for many stimuli, but not in this case. The sharp, coplanar edges dividing each of the five squares would be classified as reflectance edges, making it clear that the illuminance border must be the occluding edge forming the outer boundary of the five squares. Given the 30:1 luminance range within this illuminance boundary, the five squares would be perceived to range from black to white (under any anchoring rule). This prediction is far different from the obtained results.

The errors that occur in our staircase Gelb experiment make an especially knotty problem for intrinsic image models. One can make a provision, such as a small edge classification failure (see Gilchrist, 1988) that accounts for the erroneous lightening that occurs in a spotlight. But there is no obvious way to explain the gradient of error shown by the five squares.

12.2. Problem of area effects.

The finding that perceived lightness depends on the size of a surface does not fit comfortably within the intrinsic image approach. The intrinsic image models are models of veridical perception, based on the concept of inverse optics. For example, if reflectance and illuminance changes are confounded in the optics of the retinal image, then they need to be parsed by the visual system in an inverse manner. But the area effect cannot be viewed in this way. It is not the case that the reflectance of a surface in the physical world changes as its size changes. Thus there is no obvious way in which the area effect can be seen to counter an environmental challenge to constancy.

12.3. Problem of articulation effects.

The same argument can be made with respect to the strong effects of articulation that we have found. If surface reflectance is recovered by classifying and integrating edges, there is no reason why the number of papers within a coplanar group should have such a powerful effect on perceived lightness.

12.4. Problem of errors.

One could say that the intrinsic image models are specialized to account for veridicality while the new model is specialized to account for error, but this would be misleading. The intrinsic image models do not account properly for veridicality because they overstate it. On the other hand, the new model would not be able to make such a precise accounting of lightness errors if it were not, at the same time, able to account for veridicality, to the degree that it exists. This is a serious advantage of the anchoring model.

12.5. A response paradox.

For some years now our work has been dogged by a frustrating paradox concerning observer lightness reports. A piece of gray paper is arranged so that half of it is brightly illuminated while the other half is dimly illuminated. The two halves are divided by a blurred illuminance edge. Observers correctly report that it is an illuminance edge and that the surface lightness is the same on both sides of the edge. Yet when they make a match from a Munsell chart they always give a higher value to the lighted side than to the shadowed side. Although the discrepancy can be exacerbated when the task is misunderstood by a naive observer, the paradox runs deeper. Sophisticated observers (including the authors) make the same contradictory reports.

This paradox is a direct affront to the intrinsic image models, according to which edge classification and perceived surface lightness are tightly coupled. If surface lightness is the product of edge classification, how can such an outcome occur? But if edge classification and surface lightness are treated as parallel processes, not serial, as in the new model, such an outcome is possible.

13. New Model vs. Intrinsic Image Models

The model we are proposing has a flavor quite different from the ratiomorphic flavor of the modern computational and intrinsic image models. The emphasis on frameworks is characteristic of Gestalt theory, as is the emphasis on global relationships. Belongingness, or "appurtenance," as Koffka called it, is uniquely associated with gestalt theory. Notions of weighted average and compromise that are central to the new model are alien to the intrinsic image models, suggesting as they do perceptual forces more than the machine-like precision of edge integration or the all-or-none quality of edge classification.

The functional units into which the image is decomposed seem fundamentally different in the old and new accounts. Where the intrinsic image models picture an image composed of overlapping layers, our proposed model pictures an image composed of nested frameworks. And although the lightness computation within these frameworks probably involves both edge extraction and edge integration, there is no place for edge classification in the new model.

13.1. Classified edge integration

The absence of edge classification in the new model doesn’t mean that the visual system cannot classify edges (as noted in Section 12.5). But it does mean that the lightness system does not carry out what we call a classified integration. The process of edge integration originally proposed by Land and McCann (1971), for example, could be called an unclassified integration. That is, the luminance relationship between two widely separated regions in the retinal image is determined by integrating every edge or contour encountered along an integration path leading from one surface to the other, without regard to whether each edge represents a change in reflectance or in illuminance. But Gilchrist’s (1979) intrinsic image model specifically asserts that the visual system is able to perform a classified edge integration, an integration of only the reflectance edges between two regions, ignoring or discounting any relatively sharp illuminance borders that occur in the path of integration.

13.2. A Critical Test

Arroyo, Annan, & Gilchrist (1995) have recently conducted an experiment that pits our new model against the earlier model with its classified integration. We were surprised by the results, which appear to decisively favor the new model. The experiment is illustrated in Figure 34.

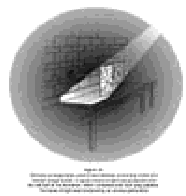

13.2.1. Method and results: New model wins.

A large, horizontally rectangular Mondrian containing about 40 gray papers was suspended in midair in a normally-illuminated laboratory room. Choice of gray papers was constrained so that only dark gray papers (between black and middle gray) were permitted on the left half of the Mondrian. The right half contained the full range of grays from black to white. A bright rectangular beam of light from an ellipsoidal spotlight was projected onto the left half of the Mondrian so that only dark gray papers were illuminated; no part of any paper lighter than middle gray was allowed to lie within the illumination.

The spotlight was not hidden in any sense. The border of the spotlight that fell across the midline of the Mondrian did not coincide with any of the reflectance borders. Indeed the illuminance border contained an obvious penumbra, and the bright illumination was readily seen by all observers.

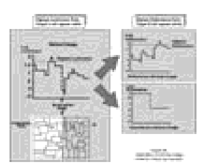

But the display constitutes an intriguing critical test, shown in schematic form in Figure 35. The question is, which target patch will appear white? Target patch A, a middle gray square in the spotlight was not only the highest luminance in its local framework but also the highest luminance in the visual field, or global framework. Thus, according to our new model, this paper should appear completely white. Target B, a white patch in the lower illumination, was the highest luminance, and thus white, in its local framework but light gray in the global framework. Given the relatively high articulation of its local framework, B should appear very light gray, almost white.

But according to Gilchrist's earlier model (1979) the visual system has the ability to perform a classified integration between the gray paper and a white paper on the non-illuminated right side of the Mondrian. An integration of all the edges, except the illuminance edge, along a path from target A to target B, a white paper outside the spotlight, should reveal that the reflectance value of target B is five times higher than that of target A. Thus, by the intrinsic image model, the visual system should be able to use the white paper as the anchor for the gray paper, even though the luminance of the gray paper is higher.

To summarize the predictions, according to the anchoring model target A should appear fully white and target B, slightly off white. According to the intrinsic image approach, target A should appear middle gray, target B, white.

The results we obtained are consistent with the new model. Target A appeared completely white. Target B appeared white as well, possibly a bit darker than target A, but the difference was not statistically significant.

13.2.2. Low articulation replication.

According to the anchoring model the difference can be increased by weakening the local framework. Annan, et al (1997) repeated the experiment under the minimum conditions of articulation that would still allow a critical test, as illustrated in Figure 36, and obtained just the results predicted. The white paper outside the spotlight was perceived as significantly darker than the gray paper in the spotlight.

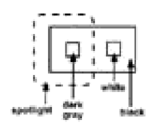

Figure 36.
Low articulation replication of critical test.

13.2.3. Review of 1983 data.

The results of our two critical tests deal a severe blow to the concept of classified edge integration. What then should we conclude about the evidence previously offered in support of this idea?

The data presented by Gilchrist, Delman, and Jacobsen (1983) and by Gilchrist (1988) were far more consistent with classified edge integration than with other models current at that time. But those data turn out to be even more consistent with our new anchoring model. Remember that the intrinsic image model predicts complete constancy, yet substantial failures of constancy appeared in the data. The anchoring model predicts those constancy failures.

14. Evaluating The Model

How should this model be critically evaluated? First, the model is undoubtedly wrong in some of its particulars. It should be regarded as a model still under construction, and we welcome contributions to its further development. Yet there must be a great deal of validity in the basic approach because no other model of lightness perception has been able to account for such a wide range of errors.

14.1. Range of application.

The greatest strength of this model is its wide range of applicability. We have surveyed an impressive array of empirical findings on lightness perception that appear quite consistent with the model.

14.2. Prediction of errors.

Our proposed model is the first model to make an explicit attempt to account for failures of lightness constancy in various situations. By comparison, the intrinsic image models predict no errors. Contrast models based on lateral inhibition predict many errors, but there is a mismatch between the errors they predict and the actual errors that occur.

14.3. Unification of constancy failures.

For some years now we have been searching for a single formula that can simultaneously account for both illumination-dependent constancy failures and background-dependent constancy failures. Historically theories of lightness perception have attempted a unification of illumination-independent constancy and background-dependent failures of constancy (see Gilchrist, 1994). Several recent attempts have been made to account for both types of constancy failure. The attempt by Gilchrist (1988) has been evaluated in Section 7.

Ross and Pessoa (in press) have presented a model in which contrast is selectively reduced at context boundaries. This approach predicts both the basic illumination-dependent and background-dependent constancy failures, at least qualitatively, and shows promise for a wider application. But we believe our proposed model provides the most widely applicable account to date.

14.4. Rigor and concreteness.

Several concepts in this model, like belongingness, framework, and strength of framework, obviously require greater clarification. Gestalt theory has long been criticized, indeed often dismissed, for a lack of rigor in its key concepts. To this charge, however, we make two responses. First, rigor should be a goal, not a fetish. Second, the alternative models are not nearly as rigorous as they seem.

We believe that we have shown a strong qualitative relationship between lightness and belongingness, intuitively defined. As in the case of the proverbial drunk, when you lose a quarter you have to look for it in the area where it was lost. There is no point in searching for it somewhere else just because the light is better. We do not claim to have offered a finished theory; we claim only that we have shown that a serious research program involving concepts of belongingness and anchoring is likely to pay rich dividends. Henneman wrote in 1935 (p. 23): "’Articulation of the visual field’ has become recognized as an essential condition for the appearance of ’constancy’ phenomena, though this rather vague term is badly in need of clearer definition and explanation." But, in the actual event, what articulation got was not a better definition, it got forgotten.

The problem of perceptual structure may be difficult but it cannot be avoided. The field of lightness perception has been plagued with attempts to account for lightness with structure-blind mechanisms. In the rush to explain perception at the physiological level as early as possible, there is a reluctance to grasp the nettle of perceptual structure.

This model represents a choice between two broad approaches to theory building. At one pole of this choice, one can begin with a domain that is sufficiently limited as to allow it to be modeled in relatively precise mathematical terms. Then one attempts to broaden the domain to which the model can be applied, without loss of rigor. At the alternative pole, one takes a wide swath of phenomena and tries to characterize, perhaps in fairly broad terms, the nature of these phenomena. Then one attempts to refine that characterization, without loss of scope. Our model is the product of the latter approach. We are defending vagueness and imprecision only as a temporary price to pay for progress.

Several models of lightness or brightness (Chubb, Sperling, & Solomon, 1989; Cornsweet, 1970; Grossberg & Todorovic', 1988; Heinemann & Chase, 1995; Jameson and Hurvich, 1964) are more formalized than the model we have presented, in the sense that they are couched in mathematical language. Yet none of these models contains a clear anchoring rule, as far as we can see. Thus, if our argument is accepted that without an anchoring rule (either explicit or implicit in the model) no model can predict specific gray shades, one must consider carefully what kinds of rigor are purchased by formalization and what kinds are not.

Neither mathematics nor physiology provides any guarantee of rigor. Mathematical language seems to confer a cloak of rigor on some models even when no one can specify exactly how the model could be falsified. Lateral inhibition is a concrete concept in that it is based on a physical process. And yet if a bright outer annulus is added to the standard disk/annulus display, there is no agreement as to whether this should depress the lightness of the disk by adding further inhibition, or enhance it through disinhibition.

The anchoring model we offer attempts to predict the specific Munsell values that will be perceived in a given image. No other theory of lightness perception makes such an attempt, even for simple images. Moreover, we can point to several steps toward greater rigor that we have already taken.

We believe that defining a framework in terms of belongingness does not merely shift the problem from one of defining framework to one of defining belongingness. Much has been known about belongingness for many years and more recently we have seen a growing accumulation of findings. Koffka (1935, p. 246) emphasizes the importance of coplanarity for belongingness in lightness perception, and we already have important findings on this topic (Adelson, 1993; Buckley, Frisby, & Freeman, 1994; Gilchrist, 1980; Knill & Kersten, 1991; Schirillo, Reeves, & Arend, 1990). Wishart has shown that the belongingness between two patches is a graded function of the angle between the planes on which they lie.

Articulation, field size, and T-junctions have been shown to be effective grouping factors across a wide range of displays. Adelson (1993) has produced several provocative new lightness illusions composed like mosaics, showing that these illusions critically depend on several types of junctions: Y-junctions, Y-junctions, and ratio-invariant X-junctions. A systematic investigation of the role of these junctions in forces of belongingness appears likely to bear fruit. Watanabe and Cavanagh (1993), Todorovic' (1996), and Anderson (1997) have already made important contributions here as well.

14.5. The role of perceived illumination.

The new model, in its current form, makes no reference to the perception of the illumination, and this may represent the biggest gap in the model. There is abundant evidence that perceived surface lightness and perceived illumination are intimately related. The classic hypothesis regarding the relationship between lightness and perceived illumination is called the albedo hypothesis or lightness-illumination invariance hypothesis (Kozaki, 1965; Kozaki & Noguchi, 1976; Logvinenko & Menshikova, 1994; Noguchi & Kozaki, 1985; Oyama, 1968). It holds that there is a unique coupling between perceived lightness and perceived illumination.

But although the model as it now stands does not specify the role of perceived illumination, there are some tantalizing findings to suggest a close connection between perceived illumination level and the anchoring of surface lightness. Several investigators have found that the perceived level of illumination is closely associated with the highest luminance in a framework (Beck, 1959, 1961), or with highest luminance and the largest area (Kozaki, 1973; Noguchi and Masuda, 1971).

14.6. Choice of scaling rule for the global framework.

A serious gap in the new model occurs because the global framework often contains a luminance range much greater than that of the black/white scale. Thus there is no obvious way that lightness values should be assigned within the global framework, although there are several possibilities. (1) The ratio principle (1:1 scaling) is the general default rule. But starting with the highest luminance and scaling all lower luminances according to the ratio rule means that all target surfaces more than 30 times darker than the maximum will be assigned the same value of black. (2) A projective scaling rule might be used. The entire dynamic range of the global framework can be scaled down proportionately to fit the 30:1 range of the black/white scale. Notice that this is equivalent to using a highest luminance rule (=white) and a lowest luminance rule (=black) at the same time. (3) A third possibility, and in our view the most likely, is that relative area plays a role in the scaling. For example, a 5:1 range of luminance values that occupy very little area within the retinal image might be mapped onto a range of lightness values much less than 5:1. Only a series of experiments can sort out the appropriate scaling rule.

15. Debt to the Early Literature

The model we are proposing owes a great deal to the early literature in lightness perception. The first three decades in this century saw a vigorous development in theory and research on surface lightness perception. Concepts like field size, articulation, and belongingness that are proving increasingly essential to an understanding of how surface lightness is processed were standard topics in those days, as can be revealed by the following list of quotations:

Burzlaff (1931, p. 25): "Katz has often emphasized that for the highest possible degree of color constancy it is very important that the field of view appear filled with numerous objects readily distinguishable from each other... The more homogeneously the visual field is filled, the more the phenomena of colour-constancy recede..."

MacLeod (1932, p32): "When, for instance a section of one’s surroundings is seen as shadowed it is of consequence whether or not the shadow occupies a large visual angle and whether the shadowed sector includes a variety of objects."

Gelb (1929, p201 in Ellis): "...colour constancy was manifested only when an adequate and uninterrupted survey of the existing illumination was permitted. Anything hindering such a survey (e.g. the reduction screen, minimal observation time, etc.) either destroyed or reduced the phenomenon. One of the most important conditions is that the visual angle be large and the field richly articulated."

Katona (1935, p61): "I have found that constancy effects are mainly furthered by enrichment of the perception through better organization, more contours, more form and object-characters, movement, etc."

Henneman (1935, p. 52): "Apparently the more complicated the constitution of the field of observation in the sense of presenting a number of surfaces of different reflectivity, regions of different illumination, and objects set in different planes (tri-dimensionality), the more in accord with its actual albedo-color will a test object be perceived."

But the historical continuity was broken by the contrast theorists, as noted in Section 1.4. Consequently several generations of younger investigators have grown up knowing little of these early ideas. The success of our proposed model underscores again the importance of citing all the relevant literature. In part, our work is an attempt to heal this historical breach.

15.1. Koffka

Overall, our proposed model owes the most to Koffka. In a discussion of color constancy, Koffka (1932) introduces the anchoring problem by invoking *the principle of the shift of level* which he notes had been advanced by Jaensch (1921) as well. Proposing what he calls an invariance theorem, Koffka writes (p. 335): "If two parts of the retina are differently stimulated, no constant relationship will exist between each part of the phenomenal field and its local stimulation, but under certain conditions there will be a constant relationship between the *gradient* in the phenomenal field and the stimulus *difference.* I.e. the two field parts may, under different conditions, look very differently coloured, but their relation one to the other or the phenomenal 'gap' between them will be the same if the stimulus difference is kept constant. The condition mentioned above is that the two parts of the field belong to the same level." This concept may be easier to understand in reference to Figure 2.

In the same passage, Koffka cites an analogous example concerning orientation. He describes "a public building on a wide lawn that slants slightly toward the lake." The building appears tilted away from the lake because the perceived horizontal anchor is tilted away from the objective horizontal by the wide lawn. Here we see the relevance of the concepts of anchoring and field size for the perception of orientation.

Enormous credit must go to Koffka and the Gestalt theorists for their emphasis on perceptual organization. In particular Koffka stressed the role of belongingness, which he called appurtenance, as the key to organization. Indeed the idea, so central to our proposed model, that belongingness is a matter of degree, rather than all-or-none, can be seen clearly in this claim by Koffka (1935, p. 246): "a field part *x* is determined in its appearance by its 'appurtenance' to other field parts. The more *x* belongs to the field part *y*, the more will its *whiteness* be determined by the gradient *xy*, and the less it belongs to the part *z*, the less will its whiteness depend on the gradient *xz*.." Koffka goes on to identify coplanarity as the most important factor producing belongingness: "Which field parts belong together, and how strong the degree of this belonging together is, depends upon factors of space organization. Clearly two parts at the same apparent distance will, *ceteris paribus*, belong more closely together than field parts organized in different planes."

Recognition must be given to a more recent Gestalt theorist who worked with Wertheimer early in her career. Dorothy Dinnerstein (1965) presented perceptual organization in terms of interacting structures, structures which in turn can be more or less strong. The impact of her ideas can be seen in the model we have presented.

15.2. Compromise and intelligence

The notion that errors might reflect the resolution of a tension between two poles is not new. Woodworth (1938, p.605) observed that judgments in constancy experiments "usually lie between two extremes, one conforming to the stimulus and the other conforming to the object". Both the Brunswik ratio and the Thouless ratio were created to measure just where the perceptual judgment lies between these two poles. Thouless (1931) spoke of "phenomenal regression to the real object." Rock (1975) attributed constancy failures to unwanted intrusions of the "proximal mode" of perception into the more object-oriented "constancy mode." Gilchrist (1988) hypothesized a competition between two processes, one that seeks to equate luminances and one that seeks to equate luminance ratios.

The issue is just how to define the two poles. Each of the above definitions is plagued with either serious logical problems or limited range of application. Empirically, however, the local/global polarity seems to successfully capture the notion of errors as products of a compromise, and to provide a wide applicability. Sedgwick and Nicholls (1993) have proposed a closely related idea in the spatial domain, viewing spatial errors as the product of cross talk between the information specifying the surface of the picture and the information specifying the three-dimensional layout of the scene. More precisely, the perception of the three-dimensional layout of the world is influenced by crosstalk from the projective relations in the optic array. At the same time, when people are instructed to report on their perception of projection relations in the optic array their reports are influenced by crosstalk from information specifying the three-dimensional layout of the scene.

Our proposed model presents a picture of visual functioning that, while far more intelligent than contrast models based on lateral inhibition, seems far less intelligent than the ratiomorphic approach that characterizes the intrinsic image models. The model implies that surface lightness is not based on a complete representation of the visual environment. As Adelson (personal communication) puts it: "processing is more powerful than low-level filtering, yet simpler than full 3-D surface description." The evidence points to a robust system that appears to accept errors even where veridical perception would seem to be possible, apparently in exchange for restricting the size of the maximum errors over a wide range of environmental challenges.

15.3. The aperture problem for lightness

Exactly what constitutes intelligent perception may not be clear for many images, even complex images. For example, imagine a richly decorated, brightly illuminated living room. There is an opening in one wall through which you can see into a second, presumably more dimly illuminated room. But you can see only two adjacent surfaces, call them targets, in the second room. How does your visual system compute the perceived gray value for those two targets? The system could anchor the targets using the bright, visually surrounding framework of the first room, by which the they would be seen as two shades of dark gray. But that computation would produce the correct values only if the illumination in the second room were equal to that in the first, and this is not known. Alternatively the system could treat that piece of the image within the aperture as a world unto itself, anchoring each of the targets relative to the other, by which they would be seen as white and light gray. This computation would produce the correct values only if the highest luminance within the aperture is really a white surface, and this is not known.

Thus the system is caught between and rock and a hard place. What is the intelligent thing to do? Making a compromise between these two alternative computations is an intelligent response, especially when the weighting in the compromise depends on factors like the number of different surfaces seen through the aperture.

In fact most of the images to which we are exposed daily contain the kind of "pockets of ambiguity" exemplified by the aperture lightness problem. Planes, apertures, and regions of special illumination that contain less than the full range of gray shades are probably the rule, rather than the exception, even in complex images. Gibson (1966) argues that ambiguity is characteristic of stimuli presented under reduced laboratory conditions, and that, in ecologically valid optic arrays, the complexity drives out the ambiguity. There is truth in Gibson’s observation, but in the case of the hole in the living room wall, increasing the optic complexity in the living room fails to reduce the ambiguity of lightness values for surfaces seen in the aperture, showing that complexity does not necessarily drive out the ambiguity. Given this reality, there seems to be little option but that a set of rules of anchoring are programmed into the system that computes surface lightness values.

REFERENCES

Adelson, E. H. (1993). Perceptual organization and the judgment of brightness. Science, 262, 2042-2044.

Adelson, E. H., & Pentland, A. P. (1990). The perception of shading and reflectance (Vision and Modeling Technical Report 140): MIT Media Laboratory.

Agostini, T., & Bruno, N. (1996). Lightness contrast in CRT and paper-and-illuminant displays. Perception & Psychophysics, 58(2), 250-258.

Agostini, T., & Galmonte, A. (1997). A new effect of gradient on lightness. Investigative Opthalmology and Visual Science, 38(4), S895.

Agostini, T., & Proffitt, D. R. (1993). Perceptual organization evokes simultaneous lightness contrast. Perception, 22(3), 263-272.

Anderson, B. (1997). A theory of illusory lightness and transparency in monocular and binocular images: The role of contour junctions. Perception, 26, 419-453.

Annan, V., Economou, E., Bonato, F., & Gilchrist, A. (1996). A paradox in surface lightness perception. Investigative Opthalmology and Visual Science, 38(4), S895.

Arend, L. (1994). Surface Colors, Illumination, and Surface Geometry: Intrinsic-Image Models of Human Color Perception. In A. Gilchrist (Ed.), Lightness, Brightness, and Transparency, (pp. 159-213). Hillsdale: Erlbaum.

Arend, L., & Goldstein, R. (1987). Simultaneous constancy, lightness and brightness. Journal of the Optical Society of America A, 4, 2281-2285.

Arend, L. E., Buehler, J. N., & Lockhead, G. R. (1971). Difference information in brightness perception. Perception and Psychophysics, 9, 367-370.

Arend, L. E., & Spehar, B. (1993). Lightness, brightness and brightness contrast. II. Reflectance variation. Perception and Psychophysics, 54(10), 457-468.

Arroyo, R., Annan, V., & Gilchrist, A. (1995). Which appears white, the highest luminance or the highest reflectance? Investigative Opthalmology and Visual Science, 36(4), S469.

Barrow, H. G., & Tenenbaum, J. (1978). Recovering intrinsic scene characteristics from images. In A. R. Hanson & E. M. Riseman (Eds.), Computer Vision Systems, (pp. 3-26). Orlando: Academic Press.

Beck, J. (1959). Stimulus correlates for the judged illumination of a surface. Journal of Experimental Psychology, 58(4), 267-274.

Beck, J. (1961). Judgments of surface illumination and lightness. Journal of Experimental Psychology, 61, 368-375.

Beck, J. (1972). Surface Color Perception: Cornell Univ. Press.

Benary, W. (1924). Beobachtungen zu einem Experiment uber Helligkeitskontrast. Psychologische Forschung, *5*, 131-142.

Bergstr_m, S. S. (1977). Common and relative components of reflected light as information about the illumination, colour, and three-dimensional form of objects. Scandinavian Journal of Psychology, 18(3), 180-186.

Bezold, W. v. (1874). Die farbenlehrer im Hinblick auf Kunst und Kuntsgewerbe (S. R. Koehler, Trans.). Brunswick: Westermann.

Bonato, F., & Cataliotti, J. (submitted). Figure/ground relations and perceived area affect the luminosity threshold. .

Bonato, F., & Gilchrist, A. (in press). Luminosity threshold depends on perceived area of a surface. .

Bonato, F., & Gilchrist, A. L. (1994). The perception of luminosity on different backgrounds and in different illuminations. Perception, 23, 991-1006.

Brown, R., & MacLeod, D. (1992). Saturation and color constancy. Advances in Color Vision Technical Digest (Optical Society of America, Washington, D. C., 1992), 4, 110-111.

Brown, R. O. (1994). The World is not Grey. Investigative Ophthalmology and Visual Science, 35(4), 2165.

Bruno, N. (1992). Lightness, equivalent backgrounds, and the spatial integration of luminance. Perception supplement(21), 80.

Bruno, N. (1994). Failures of lightness constancy, edge integration, and local edge enhancement. Vision Research, 34(17), 2205-2214.

Bruno, N., Bernardis, P., & Schirillo, J. (In press). Lightness, equivalent backgrounds, and anchoring. Perception & Psychophysics.

Buchsbaum, G. (1980). A spatial processor model for object color perception. Journal of the Franklin Institute, 310, 1-26.

Buckley, D., Frisby, J., & Freeman, J. (1994). Lightness perception can be affected by surface curvature from stereopsis. Perception, 23, 869-881.

Burgh, P., & Grindley, G. C. (1962). Size of test patch and simultaneous contrast. Quarterly Journal of Experimental Psychology, 14, 89-93.

Burzlaff, W. (1931). Methodologische Beitrage zum Problem der Farbenkonstanz. Zeitschrift für Psychologie, 119, 117-235.

Cataliotti, J., & Gilchrist, A. L. (1995). Local and global processes in lightness perception. Perception and Psychophysics, 57(2), 125-135.

Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. Proceedings of the National Academy of Sciences USA, 86, 9631-9635.

Cole, R. E., & Diamond, A. L. (1971). Amount of surround and test inducing separation in simultaneous brightness contrast. Perception and Psychophysics, 9, 125-128.

Cornsweet, T. N. (1970). Visual Perception. New York: Academic Press.

Cramer, T. (1923). ?ber die Beziehung des Zwischenmediums zu den Transformations- und Konstrasterscheinungen (The effect of intervening media on transformation- and contrast-appearances). Zeitschrift für Sinnesphysiologie, 54(2), 215-242.

De Valois, R. L., & De Valois, K. K. (1988). Spatial Vision. New York: Oxford University Press.

Diamond, A. (1953). Foveal simultaneous brightness contrast as a function of inducing- and test-field luminances. Journal of Experimental Psychology, 45, 304-314.

Diamond, A. (1955). Foveal simultaneous brightness contrast as a function of inducing- and test-field luminances. Journal of Experimental Psychology, 50, 144-152.

Diamond, A. (1960). A theory of depression and enhancement in the brightness response. Psychological Review, 67(3), 168-199.

Diamond, A. (1962). Brightness of a field as a function of its area. Journal of the Optical Society of America, 52, 700-706.

Dinnerstein, D. (1965). Previous and concurrent visual experience as determinants of phenomenal shape. American Journal of Psychology, 78, 235-242.

Dunn, B., & Leibowitz, H. (1961). The effect of separation between test and inducing fields on brightness constancy. Journal of Experimental Psychology, 61(6), 505-507.

Evans, R. (1959). Fluorescence and gray content of surface colors. Journal of the Optical Society, 49, 1049.

Evans, R. M. (1964). Variables of perceived color. Journal of the Optical Society of America, 54, 1467-1474.

Evans, R. M. (1974). The Perception of Color. New York: Wiley.

Flock, H. R., & Noguchi, K. (1970). An experimental test of Jameson and Hurvich's theory of brightness contrast. Perception and Psychophysics, 8, 129-136.

Freeman, R. B. (1967). Contrast interpretation of brightness constancy. Psychological Bulletin, 67, 165-187.

Fry, G. A., & Alpern, M. (1953). The effect of a peripheral glare source upon the apparent brightness of an object. Journal of the Optical Society of America 43, 189-195.

Gelb, A. (1929). Die "Farbenkonstanz" der Sehdinge. In W. A. von Bethe (Ed.), Handbuch der Normal und Pathologische Psychologie, (pp. 594-678).

Gerbino, W. (1994). Achromatic Transparency. In A. Gilchrist (Ed.), Lightness, Brightness, and Transparency, (pp. 215-255). Hillsdale: Erlbaum.

Gibbs, T., & Lawson, R. B. (1974). Simultaneous brightness contrast in stereoscopic space. Vision Research, 14, 983-987.

Gibson, J. J. (1966). The Senses Considered as Perceptual Systems. Boston: Houghton Mifflin.

Gilchrist, A. (1979). The perception of surface blacks and whites. Scientific American, 240, 112-123.

Gilchrist, A. (1988). Lightness contrast and failures of constancy: a common explanation. Perception and Psychophysics, 43(5), 415-424.

Gilchrist, A. (1994). Absolute Versus Relative Theories of Lightness Perception. In A. Gilchrist (Ed.), Lightness, Brightness, and Transparency, (pp. 1-33). Hillsdale: Erlbaum.

Gilchrist, A. (1996). The deeper lesson of Alhazen. Perception, 25, 1133-1136.

Gilchrist, A., & Bonato, F. (1995). Anchoring of lightness values in center/surround displays. Journal of Experimental Psychology: Human Perception and Performance, 21(6), 1427-1440.

Gilchrist, A., & Cataliotti, J. (1994). Anchoring of surface lightness with multiple illumination levels. Investigative Opthalmology and Visual Science, 35(4), S2165.

Gilchrist, A. L. (1977). Perceived lightness depends on perceived spatial arrangement. Science, 195(4274), 185-187.

Gilchrist, A. L. (1980). When does perceived lightness depend on perceived spatial arrangement? Perception and Psychophysics, 28(6), 527-538.

Gilchrist, A. L., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. Perception and Psychophysics, 33(5), 425-436.

Gilchrist, A. L., & Jacobsen, A. (1983). Lightness constancy through a veiling luminance. Journal of Experimental Psychology: Human Perception and Performance, 9(6), 936-944.

Gilchrist, A. L., & Jacobsen, A. (1988). Qualitative relationships are decisive. Perception and Psychophysics, Notes and Comment, 45 (1), 92-94.

Gogel, W. C., & Mershon, D. H. (1969). Depth adjacency in simultaneous contrast. Perception and Psychophysics, 5(1), 13-17.

Grossberg, S., & Todorovic, D. (1988). Neural dynamics of 1-d and 2-d brightness perception: A unified model of classical and recent phenomena. Perception and Psychophysics, 43, 241-277.

Haimson, B. R. (1974). The response criterion, the stimulus configuration, and the relationship between brightness contrast and brightness constancy. Perception and Psychophysics, 16 (2), 347-354.

Heinemann, E., & Chase, S. (1995). A quantitative model for simultaneous brightness induction. Vision Research, 35(14), 2007-2020.

Heinemann, E. G. (1955). Simultaneous brightness induction as a function of inducing-and test-field luminances. Journal of Experimental Psychology, 50(2), 89-96.

Heinemann, E. G. (1971). Simultaneous Brightness Induction. In D. Jameson & L. Hurvich (Eds.), Handbook of Sensory Physiology, .

Heinemann, E. G. (1988). Brightness contrast, brightness constancy, and the ratio principle. Perception and Psychophysics, 45, 89-91.

Helmholtz, H. v. (1866). Helmholtz’s treatise on physiological optics. New York: Optical Society of America.

Helson, H. (1943). Some factors and implications of color constancy. Journal of the Optical Society of America, 33(10), 555-567.

Helson, H. (1963). Studies of anomalous contrast and assimilation. Journal of the Optical

Society of America, 53, 179-184.

Helson, H. (1964). Adaptation-Level Theory. New York: Harper & Row.

Helson, H., & Joy, V. (1962). Domains of lightness, assimilation and contrast effects in vision. Psychologische Beitraege, 6, 405-415.

Helson, H., & Rohles, F. H. (1959). A quantitative study of reversal of classical lightness contrast. American Journal of Psychology, 72, 530-538.

Henneman, R. H. (1935). A photometric study of the perception of object color. Archives of Psychology, No. 179 (No. 179), 5-89.

Hering, E. (1874, 1964). Outlines of a Theory of the Light Sense. (L.M. Hurvich & D. Jameson, Trans.). Cambridge, Mass.: Harvard University Press.

Horeman, H. W. (1963). Inductive brightness depression as influenced by configurational conditions. Vision Research, 3, 121-130.

Horeman, H. W. (1965). Relations between brightness and luminance under induction. Vision Research, 5, 331-340.

Horn, B. K. P. (1977). Understanding image intensities. Artificial Intelligence, 21(11), 201-231.

Horn, B. K. P. (1986). Robot Vision. Cambridge: MIT Press.

Hurlbert, A. (1986). Formal connections between lightness algorithms. Journal of the Optical Society of America A, 3, 1684-1693.

Hurvich, L., & Jameson, D. (1966). The Perception of Brightness and Darkness. Boston: Allyn and Bacon.

Jacobsen, A., & Gilchrist, A. (1988). Hess and Pretori revisited: Resolution of some old contradictions. Perception and Psychophysics, 43, 7-14.

Jacobsen, A., & Gilchrist, A. (1988). Hess and Pretori revisited: Resolution of some old contradictions. Perception and Psychophysics, 43, 7-14.

Jacobsen, A., & Gilchrist, A. (1988). The ratio principle holds over a million-to-one range of illumination. Perception and Psychophysics, 43, 1-6.

Jacobsen, A., & Gilchrist, A. (1988). The ratio principle holds over a million-to-one range of illumination. Perception and Psychophysics, 43, 1-6.

Jaensch, E. (1921). ?ber den Farbenkontrast und die sogenannte Berucksichtigung der Beleuchtung. Zeitschrift für Sinnesphysiologie, 52, 165-180.

Jameson, D., & Hurvich, L. M. (1961). Complexities of perceived brightness. Science, 133, 174-179.

Jameson, D., & Hurvich, L. M. (1964). Theory of brightness and color contrast in human vision. Vision Research, 4, 135-154.

Judd, D. (1966). Basic correlates of the visual stimulus. In S. S. Stevens (Ed.), Handbook of Experimental Psychology, (pp. 811-867). New York: Wiley.

Kanizsa, G. (1979). Organization in Vision. New York: Praeger.

Kardos, L. (1934). Ding und Schatten. Zeitschrift für Psychologie, Erg. bd 23.

Kardos, L. (1935). Versuch einer mathematischen analyse von gesetzen des farbensehens. nahere bestimmung des funktionalen verhaltnisses zwishen farbenerlebnis und reizgesamtheit. Zeitschrift für Sinnesphysiologie, 66, 182-212.

Katona, G. (1929). Zur Analyse der Helligkeitskonstanz. Psychologische Forschung, 12, 94-126.

Katona, G. (1935). Color constrast and color constancy. Journal of Experimental Psychology, 18, 49-63.

Katz, D. (1911). Die Erscheingsweisen der Farben und ihre Beeinflussung durch die individuelle Erfahrung. Zeitschrift für Psychologie, 7.

Katz, D. (1935). The World of Colour. London: Kegan Paul, Trench,Trubner & Co.

Knill, D., & Kersten, D. (1991). Apparent surface curvature affects lightness perception. Nature, 351(May), 228-230.

Koffka, K. (1932). Some remarks on the theory of colour constancy. Psychologische Forschung, 16, 329-354.

Koffka, K. (1935). Principles of Gestalt Psychology. New York: Harcourt, Brace, and World.

Kozaki, A. (1963). A further study in the relationship between brightness constancy and contrast. Japanese Psychological Research, 5, 129-136.

Kozaki, A. (1965). The effect of co-existent stimuli other than the test stimulus on brightness constancy. Japanese Psychological Research, 7, 138-147.

Kozaki, A. (1973). Perception of lightness and brightness of achromatic surface color and impression of illumination. Japanese Psychological Research, 15(4), 194-203.

Kozaki, A., & Noguchi, K. (1976). The relationship between perceived surface-lightness and perceived illumination. Psychological Research, 39(1), 1-16.

Land, E. H. (1983). Recent advances in retinex theory and some implications for cortical computations: Color vision and the natural image. Proceedings of the National Academy of Sciences USA, Physics, 80, 5163-5169.

Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. Journal of the Optical Society of America, 61, 1-11.

Laurinen, P. I., & Olzak, L. A. (1994). Spatially Remote, Frequency-Tuned Mechanism can Control Apparent Lightness. Perception, 23, 18.

Leibowitz, H. (1965). Visual Perception. New York: MacMillan.

Leibowitz, H., Mote, F. A., & Thurlow, W. R. (1953). Simultaneous contrast as a function of separation between test and inducing fields. Journal of Experimental Psychology, 46(6), 453-456.

Leibowitz, H., Myers, N. A., & Chinetti, P. (1955). The role of simultaneous contrast in brightness constancy. Journal of Experimental Psychology, 50(1), 15-18.

Li, X., & Gilchrist, A. (in press). Relative area and relative luminance combine to anchor surface lightness values. .

Lie, I. (1969). Psychophysical invariants of achromatic colour vision IV. Depth adjacency and simultaneous contrast. Scandanavian Journal of Psychology, 10, 282-286.

Logvinenko, A., & Menshikova, G. (1994). Trade-off between achromatic colour and perceived illumination as revealed by the use of pseudoscopic inversion of apparent depth. Perception, 23(9), 1007-1024.

MacLeod, R. B. (1932). An experimental investigation of brightness constancy. Archives of Psychology(No. 135), 5-102.

MacLeod, R. B. (1940). Brightness constancy in unrecognized shadows Journal of Experimental Psychology, 27(1), 1-22.

MacLeod, R. B. (1947). The effects of artificial penumbrae on the brightness of included areas, Miscellanea Psychologica Albert Michotte, (pp. 138-154). Louvain: Institut Superieur de Philosophie.

Marr, D. (1982). Vision. San Francisco: Freeman.

McCann, J. (1994,). Psychophysical Experiments in Search of Adaptation and the Gray World. Paper presented at the IS&T Annual Meeting, Rochester, NY.

McCann, J. J. (1987). Local/Global Mechanisms for Color Constancy. Die Farbe, 34, 275-283.

McCann, J. J., & Savoy, R. L. (1991). Measurements of lightness: Dependence on the position of a white in the field of view. SPIE, 1453, 402-411.

Mershon, D., & Gogel, W. (1970). Effect of stereoscopic cues on perceived whiteness. American Journal of Psychology, 83, 55-67.

Mershon, D. H. (1972). Relative contributions of depth and directional adjacency to simultaneous whiteness contrast. Vision Research, 12, 969-979.

Metelli, F. (1985). Stimulation and perception of transparency. Psychological Research, 47, 185-202.

Metelli, F., Da Pos, O., & Cavedon, A. (1985). Balanced and unbalanced, complete and partial transparency. Perception and Psychophysics, 38, 354-366.

Newson, L. J. (1958). Some principles governing changes in the apparent lightness of test surfaces isolated from their normal backgrounds. Quarterly Journal of Experimental Psychology, 10, 82-95.

Noguchi, K., & Kozaki, A. (1985). Perceptual scission of surface-lightness and illumination: An examination of the Gelb effect. Psychological Research, 47(1), 19-25.

Noguchi, K., & Masuda, N. (1971). Brightness changes in a complex field with changing illumination: A re-examination of Jameson and Hurvich's study of brightness constancy. Japanese Psychological Research, 13, 60-69.

Oyama, T. (1968). Stimulus determinants of brightness constancy and the perception of illumination. Japanese Psychological Research, 10, 146-155.

Reid, R. C., & Shapley, R. (1988). Brightness induction by local contrast and the spatial dependence of assimilation. Vision Research, 28, 115-132.

Rock, I. (1973). Orientation and Form. New York: Academic Press.

Rock, I. (1975). An Introduction to Perception. New York: Macmillan.

Rock, I. (1983). The Logic of Perception. Cambridge, Massachusetts: MIT Press.

Rock, I., & Brosgole, L. (1964). Grouping based on phenomenal proximity. Journal of Experimental Psychology, 67(6), 531-538.

Ross, W., & Pessoa, L. (in press) Contrast/filling-in model of 3-D lightness perception. Perception and Psychophysics.

Schirillo, J., & Shevell, S. (1996). Brightness contrast from inhomogeneous surrounds. Vision Research, 36(12), 1783-1796.

Schirillo, J. A., & Arend, L. (1995). Illumination change at a depth edge can reduce constancy. Perception and Psychophysics, 57(2), 225-230.

Schirillo, J. A., Reeves, A., & Arend, L. (1990). Perceived lightness, but not brightness, of achromatic surfaces depends on perceived depth information. Perception and Psychophysics, 48(1), 82-90.

Schouten, S., & Blommaert, F. (1995). Brightness constancy in a Ganzfeld environment. Perception and Psychophysics, 57 (7), 1112-1122.

Schouten, S., & Blommaert, F. (1995). Brightness indention: A novel compression mechanism in the luminance-brightness mapping. Perception and Psychophysics, 57 (7), 1023-1031.

Sedgwick, H., & Nicholls, A. (1993). Cross talk between the picture surface and the pictured scene: Effects on perceived shape". Perception, 22 (supplement), 109.

Semmelroth, C. C. (1970). Prediction of lightness and brightness on different backgrounds. Journal of the Optical Society of America, 60(12), 1685-1689.

Shapley, R., & Enroth-Cugell, C. (1984). Visual adaptation and retinal gain controls. Progess in Retinal Research, 3, 263-343.

Shapley, R., & Reid, R. C. (1985). Contrast and assimilation in the perception of brightness. Proceedings of the National Academy of Sciences of the USA, 82, 5983-5986.

Shimojo, S., & Nakayama, K. (1990). Amodal presence of partially occluded surfaces: role of invisible stimuli in apparent motion correspondence. Perception, 19, 285-299.

Spehar, B., Gilchrist, A., & Arend, L. (1995). White’s illusion and brightness induction: The critical role of luminance relations. Vision Research, 35, 2603-2614.

Stevens, J. C. (1967). Brightness inhibition re size of surround. Perception and Psychophysics, 2, 189-192.

Stewart, E. (1959). The Gelb effect. Journal of Experimental Psychology, 57, 235-242.

Takasaki, H. (1966). Lightness change of grays induced by change in reflectance of gray background. Journal of the Optical Society of America, 56(4), 504-509.

Taya, R., Ehrenstein, W., & Cavonius, C. (1995). Varying the strength of the Munker-White effect by stereoscopic viewing. Perception, 24, 685-694.

Thouless, R. H. (1931). Phenomenal regression to the "real" object. British Journal of Psychology, 22, 1-30.

Todorovic¢ , D. (1997). Lightness and junctions. Perception, 26(4), 379-394.

Torii, S., & Uemura, Y. (1965). Effects of inducing luminance and area upon the apparent brightness of a test field. Japanese Psychological Research, 7(2), 86-100.

Ullman, S. (1976). On visual detection of light sources. Biological Cybernetics, 21, 205-212.

Wallach, H. (1948). Brightness constancy and the nature of achromatic colors. Journal of Experimental Psychology, 38, 310-324.

Wallach, H. (1963). The perception of neutral colors. Scientific American, 208, 107-116.

Wallach, H. (1976). On Perception. New York: Quadrangle/The New York Times Book Co.

Watanabe, T., & Cavanagh, P. (1993). Transparent surfaces defined by implicit X junctions. Vision Research, 33, 2339-2346.

Whipple, W., R., Wallach, H., & Marshall, F. J. (1988). The effect of area, separation, and dichoptic presentation on the perception of achromatic color. Perception and Psychophysics, 43(4), 367-372.

White, M. (1981). The effect of the nature of the surround on the perceived lightness of grey bars within square-wave test gratings. Perception, 10(2), 215-230.

Whittle, P. (1992). Brightness, discriminability and the ’Crispening Effect’. Vision Research, 32(8), 1493-1507.

Whittle, P. (1994). Contrast brightness and ordinary seeing. In A. Gilchrist (Ed.), Lightness, Brightness, and Transparency, (pp. 111-158). Hillsdale: Erlbaum.

Whittle, P. (1994). The psychophysics of contrast brightness. In A. Gilchrist (Ed.), Lightness, Brightness, and Transparency, (pp. 35-110). Hillsdale: Erlbaum.

Whittle, P., & Challands, P. D. C. (1969). The effect of background luminance on the brightness of flashes. Vision Research, 9, 1095-1110.

Wishart, K., Frisby, & Buckley. (1995). The Role of Contour in a Lightness Illusion. Investigative Ophthalmology & Visual Science, 36(4), S640.

Wolff, W. (1933). Concerning the contrast-causing effect of transformed colors. Psychologische Forschung, 18, 90-97.

Wolff, W. (1934). Psychologische Forschung, 20, 159-194.

Woodworth, R. S. (1938). Experimental Psychology. New York: Holt.

Yarbus, A. L. (1967). Eye Movements and Vision. New York: Plenum Press.

Yund, E. W., & Armington, J. C. (1975). Color and brightness contrast effects as a function of spatial variables. Vision Research, 15(8-9), 917-929.

Author Notes

The authors wish to acknowledge the support of the National Science Foundation (grant numbers BNS-8909182, DBS-9222104 and SBR 95-14679).

We would like to thank the following people who commented on an earlier draft: Vidal Annan, Larry Arend, Nicola Bruno, Elias Economou, Dave Emmerich, Alesandra Galmonte, Walter Gerbino, Holger Knau, Luiz Pessoa, William Ross, Hal Sedgwick, Dejan Todorovic', and Paul Whittle.

Requests for reprints should be addressed to Alan Gilchrist, Psychology Dept., Rutgers University, Newark, NJ 07102.

Figure Captions

Figure 1. The ambiguity of luminance ratios.

Figure 2. The anchoring problem.

Figure 3. 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.

Figure 4. Evidence for the highest luminance rule, showing the actual range of grays and the perceived range in three separate experiments.

Figure 5. Actual and perceived gray shades in domes experiments (Li & Gilchrist, in press).

Figure 6. Schematic showing how lightness varies as a function of relative area in a two-field dome with constant luminance values.

Figure 7. Belongingness is increased across the stem of the T, but decreased across the top of the T in a T-junction.

Figure 8. Perceived lightness range for five squares in spotlight is compressed relative to the actual range.

Figure 9. Schematic showing theoretical account of compression.

Figure 10. Dependence of compression on configuration of targets. Mondrian configuration produces stronger local anchoring.

Figure 11. Dependence of compression on number of squares. Greater articulation (more squares) produces stronger local anchoring.

Figure 12. Dependence of compression on size of display. Larger field size produces stronger local anchoring.

Figure 13. Greater perceived area produces higher luminosity threshold (Bonato & Cataliotti, in preparation).

Figure 14. Insulation: white border produces stronger local anchoring; black border has no effect.

Figure 15. Test of contrast interpretation of borders effect. Lightness of black target shows little or no dependence on degree of contact with white border.

Figure 16. Symmetry in illumination-dependent and background-dependent failures of constancy.

Figure 17. Katz paradigm for studying illumination-independent lightness constancy.

Figure 18. Burzlaff experiments on lightness constancy and articulation.

Figure 19. Role of articulation in two experiments by Schirillo and colleagues.

Figure 20. Simultaneous lightness contrast display, showing the stronger anchoring component and the weaker scaling component.

Figure 21. Locus of the error in simultaneous lightness contrast. Veridical value is 1.30.

Figure 22. Standard version of simultaneous contrast compared with double increments version and high articulation version.

Figure 23. Benary effect: triangles appear to differ in lightness even though they have identical local neighbors.

Figure 24. White's illusion (White, 1981).

Figure 25. Checkerboard contrast (DeValois & DeValois, 1988).

Figure 26. Schematic of Agostini & Proffitt (1993) experiment. Gray disk moving with black disks appears slightly lighter than gray disk moving with white disks.

Figure 27. Wolff illusion (1934). The light and dark areas are equal in each figure. But the background regions are perceived to be larger and thus the Area Rule predicts that each region in the left hand display will appear lighter than the region of equal luminance in the right hand display.

Figure 28. Adelson's (1993) corrugated Mondrian. Target B appears lighter than target A.

Figure 29. Application of anchoring model to Wishart experiments. The whole pattern of results here is consistent with anchoring if belongingness varies with coplanarity and coplanarity is treated as a graded variable.

Figure 30. If the illusion is due to perception of target B in a shadowed plane, it should still be present. But the anchoring model predicts no illusion here because of the lack of articulation in the local framework of B.

Figure 31. Application of highest luminance rule to Heinemann (1955) data.

Figure 32. Application of area rule to data by Stevens and Diamond. Area effects occur mainly under conditions, shown as shaded regions, specified by the Area Rule.

Figure 33. The lightness of the black square in spotlight as lighter squares are added. The straight line shows that the darkening of the black square depends on the highest luminance but not on its distance from the highest luminance.

Figure 34. Stimulus arrangements used to test between anchoring model and intrinsic image model. A square beam of light was projected onto the left half of the Mondrian, which contained only dark gray patches. The beam of light was bordered by an obvious penumbra.

Figure 35. Application of intrinsic image model to critical test stimulus.

Figure 36. Low articulation replication of critical test.