Fitting the world to the mind: Transforming images to mimic perceptual adaptation

Michael A. Webster[1], Kyle McDermott[1], and George Bebis[2]

Departments of Psychology[1] and Computer Science and Electrical Engineering[2]
University of Nevada, Reno

Abstract. Visual sensitivity is constantly adjusting to the current visual context through processes of adaptation. These adaptive changes strongly affect all perceptual judgments and optimize visual coding for the specific properties of the scenes before us. As a result human observers "see" better when they are first allowed to adapt to a specific context. The basic form of the response changes resulting from adaptation have been studied extensively and many are known in broad outline. Here we consider the advantages of applying these changes to images, to simulate the processes of adaptation that normally occur within the observer. Matching images to the observer may obviate the need for some forms of perceptual learning and offers a number of potential benefits for interacting with visual displays.

1 Introduction

Visual perception is a highly dynamic process, adjusted continuously by mechanisms of adaptation that recalibrate visual coding according to the statistics of the image we are currently viewing [1]. Numerous classic examples illustrate the marked adaptability of perception and how changes in the states of adaptation profoundly alter the way the world looks. For example, after viewing a red square a gray square appears greenish; exposure to a tilted line causes a vertical line to appear tilted in the opposite direction; and after a few moments looking at the downward flow of a waterfall, the static rocks to the side appear to ooze upward. Adaptation aftereffects occur not only for simple stimulus dimensions but also for highly complex and abstract image properties. For example, the perceived configuration of a face can be strongly biased by prior exposure to a distorted face [2] or to faces drawn from different identities [3] or facial categories such as gender, ethnicity, or expression [4]. Such results suggest that the sensitivity changes underlying adaptation occur at all stages of visual coding and may in fact be an intrinsic and fundamental property of neural activity.

Adaptation is thought to confer a number of interrelated functional advantages to the observer. Many of these have been postulated based on information theory, though the extent to which they can be linked to actual improvements in behavioral performance remains to be established. Some potential benefits of adaptation include:

1. Maximizing the limited dynamic range available for visual coding. The range of response levels a neuron can reliably transmit is highly restricted and may be orders of magnitude smaller than the range of stimulus levels that are typically encountered. Thus this range must be centered around the ambient stimulus level much as a camera's settings are adjusted to match the limited operating characteristic of the film or CCD array. The clearest example of the consequent improvements for seeing is in light adaptation, in which visual sensitivity adjusts to the mean light level so that the "exposure" level remains at an appropriate level for perceiving the variations in light around the mean [5]. Similar adjustments may also be important for coding image contrasts. For example, cells in visual cortex code only a limited range of contrasts and must adapt to the ambient contrast level in scenes to avoid response saturation [6].

2. Improving visual discrimination. Adjusting to the mean stimulus level in images allows differences around the mean to be more easily distinguished. Again, we are sensitive to variations in light only around the average light level we are adapted to, and similarly in color vision, discrimination is best for stimuli near the white point and falls progressively for saturated colors [7]. Analogous processes might also underlie high-level perceptual judgments such as the "other race effect" in face perception, in which we can readily discriminate differences between faces within the ethnic group we are exposed to while faces drawn from novel groups appear similar.

3. Improving coding efficiency within mechanisms. To carry the most information a neuron's responses should be matched to the distribution of stimuli so that any given response level occurs with equal probability [8]. This means that responses should change rapidly where stimulus levels are abundant while asymptoting where stimulus levels are rare, similar to histogram equalization. It is unclear to what extent actual processes of adaptation can adjust to the specific distribution of stimulus levels and whether these adjustments differ for different stimulus dimensions. However, in color coding it is clear that adaptation independently alters responses to both the mean and the variance of the stimulus [9].

4. Improving coding efficiency across mechanisms. At least at peripheral stages, information in the visual system is thought to be represented within filters or channels that span different ranges of the stimulus dimension [10]. Thus color is coded by three cone types with different spectral sensitivities while orientation or size are represented by multiple channels each tuned to a limited range of tilts or spatial scales. By allowing each channel to adjust to the mean stimulus level it is exposed to, responses are equated across the population of mechanisms. This maximizes efficiency by allowing each channel to carry the same amount of potential information [11].

5. Removing redundancy. Some forms of adaptation might include mutual inhibition between channels whenever they respond together to the stimulus. Such processes could improve coding efficiency by decorrelating the responses across the set of channels, allowing each to carry independent information about the stimulus [12]. A possible sign of these adjustments is contingent aftereffects, in which observers adapt to the correlations between stimulus dimensions. For example, in the McCollough Effect, viewing red vertical bars alternated with green horizontal bars causes a vertical gray bar to appear greenish while horizontal bars appear reddish. The color aftereffect is thus contingent on the pairing of color with orientation (and also on the pairing of chromaticity and luminance in the image). The negative color aftereffects are consistent with the responses predicted by decorrelating the joint responses to color and luminance and tilt in the stimulus [12], but could also arise from independent sensitivity changes within mechanisms selectively tuned to different combinations of orientation and color.

6. Maintaining perceptual constancy. The preceding examples all share the potential advantage of improving the discriminative capacity of perception by maximizing the information carried by the neural response. A somewhat different goal of perception may be to maximize recognition – to allow a given visual response to faithfully represent a consistent property of the stimulus. The problem of recognition often involves extracting invariant characteristics of the object by discounting extraneous sources of variation in the visual response. The role of adaptation in solving this problem has been most thoroughly investigated in the case of color constancy, in which the goal of the observer is to recognize the same surface under different viewing conditions [13]. This problem is ill-posed because the color signal reaching the observer depends on both the surface reflectance function and the spectrum of the incident illumination. Many different processes contribute to color constancy in human vision, yet adapting to the average color in the scene can often lead to approximate constancy by removing the changes in the average color owing to the illuminant [13]. The same process may also underlie a more important form of constancy by removing variations in the observer. For example, the lens of the human eye includes an inert pigment that selectively absorbs short-wavelength light. The density of the lens pigment increases with age and thus the average light spectrum reaching the receptors is very different in younger and older observers. However, color appearance changes little with age, because adaptation continuously readjusts the visual response to discount the changes in visual sensitivity [14]. Similar adjustments may account for the constant perception of image focus despite the marked differences in visual acuity across the visual field or as spatial sensitivity varies during development and aging [15].

7. Building predictive codes. Finally, by adjusting to the current stimulus adaptation may allow the visual code to take the form of a prediction about the world [16]. The concept of predictive codes is closely related to norm-based codes in visual perception. In these the stimulus is represented by how it deviates from an average or prototype stimulus. Many visual dimensions may be represented relative to norms, including color, shape, and motion, and high-level properties such as facial identity [17]. Adaptation is thought to play a critical role in such codes by establishing and updating the norm based on the average stimulus history. This has the advantage that resources can be devoted to signaling only the errors in the prediction, which are generally the most informative features of the image. In particular, adapting to the average properties of images may increase the perceptual salience of novel image features and thus more readily allow the visual system to detect statistical outliers or "suspicious coincidences" in the environment [12].

As the foregoing list illustrates, most accounts of adaptation assume that the response changes are designed to improve visual performance. That is, an observer can see better if they are appropriately adapted to the image they are currently viewing. Normally this requires changing the sensitivity of the observer in order to match visual coding to the current environment. Thus the processes of adaptation involve "Fitting the mind to the world" [18]. However, image processing provides the potential for increasing visual performance by instead adapting the image to the observer, or “fitting the world to the mind.” To the extent that the visual transforms underlying adaptation are known, images might be processed to simulate the changes in perception that would normally occur through adaptation, obviating the need for adjustments in the observer. Here we consider some of the transforms that are likely to occur and their implications for particular visual tasks. Our aim is not to offer a specific transform, but rather to note the potential for developing algorithms that might improve human performance with visual displays.

2 Processes of adaptation

Here we describe the visual response changes that actually occur during perceptual adaptation. While we remain far from a complete understanding of these changes, certain general principles appear well established and common to many stimulus dimensions. These principles include the following factors:

1. Adjusting to the mean. Many adaptation phenomena involve changes in sensitivity driven by the average properties of the stimulus [1]. For example, light and chromatic adaptation reflect adjustments to the average luminance and chromaticity of the stimulus, while motion adaptation involves a change in the response to the average direction. Changes in sensitivity to the mean stimulus are the most ubiquitous property of visual adaptation and lead to the largest and most compelling visual aftereffects.

2. Normalization vs. repulsion. For many dimensions, adaptation to the mean causes the stimulus to appear more neutral and thus represents a renormalization of the visual code [17]. To draw again from color, viewing a red field causes the field to appear more desaturated and in the limit to appear achromatic. The currently viewed stimulus thus becomes the new norm or neutral point for color coding. In other cases adaptation is found to reduce sensitivity to the adapting stimulus without changing its appearance. A well known example of this is spatial frequency adaptation, in which subjects are adapted to a spatially varying grating of a particular frequency [10]. While the grating itself may appear reduced in contrast, its perceived size remains unaltered. However, the perceived size of gratings higher or lower than the adapting level appear biased or “repulsed” away from the adapting frequency. In general, repulsion tends to occur for stimulus dimensions like size or frequency, which do not have a unique neutral point, while normalization is more likely to occur for stimulus dimensions like color that include a clear norm.

3. Subtractive vs. multiplicative response changes. In some cases adaptation tends to reflect a simple gain change in visual sensitivity, equivalent to rescaling the response by a constant. Such scaling closely describes some of the sensitivity changes in chromatic adaptation. In other cases the adapted response is better described by subtracting a fixed constant from the response. A complete account of light and chromatic adaptation in fact includes both forms of adjustment, with multiplicative scaling of the receptor sensitivities and subtractive adjustments from spatial and temporal filtering [19]. The response changes underlying pattern adaptation may also approximate a subtractive loss in sensitivity [20], but have not been well characterized for many dimensions. The nature of the change has important implications for modeling the perceptual consequences of adaptation. For example, multiplicative changes have large effects at all contrast levels while subtractive changes become negligible at high contrasts where the visual response is already strong.

4. Adjusting to the variance. Adaptation adjusts not only to the average stimulus level but also to the range of variation around the average, or to contrast. For example, after viewing a pattern that alternates between red and green, reds and greens appear less saturated relative to other colors [1]. There have been comparatively few studies of how the visual system adapts to variance in other stimulus dimensions, but such adjustments seem necessary in order to exploit the full dynamic range of the visual mechanisms encoding these dimensions.

5. Selectivity of visual aftereffects. Visual aftereffects are selective – viewing a tilted line reduces sensitivity to lines with similar tilt, but has much less effect on the detectability or appearance of lines at more distant orientations [21]. Conversely, adapting to a particular color may lead to more uniform changes in appearance throughout color space [9]. These differences in part reflect how narrowly or broadly tuned the mechanisms are that encode the stimulus dimension. A large volume of research has focused on characterizing the number and selectivity of visual channels in different visual tasks [10]. However, this remains largely an empirical question so that it is not yet possible to specify a uniform principle that might guide the choice of filters for modeling a particular adaptation effect.

3 Formal models of adaptation

A number of computational models have now been developed to predict adaptation effects in the human visual system. These incorporate many of the basic processes outlined above and thus offer the potential to devise image transforms that simulate adaptation. Color coding is again the domain that has seen the most attention and for which current models have come closest to accurate quantitative predictions [19] [22]. As noted above, light and chromatic adaptation involve multiplicative gain changes in the cones followed by subtractive adjustments in post-receptoral processes. Models of these processes can now closely account for visual sensitivity to luminance and color under a wide variety of viewing conditions, and are at a point where they could be readily implemented in visual displays.

Models for adaptation processes at more central sites in the visual system have yet to reach the same degree of rigor, yet have nevertheless succeeded in characterizing the general form of aftereffects observed in pattern-selective adaptation. A number of these studies have combined gain control within the channel responses and decorrelation between the channel responses in order to account for the selective changes in sensitivity that follow adaptation to color contrast or orientation or motion [23-26]. However, essentially the same perceptual changes can also be modeled by instead assuming independent response changes across multiple mechanisms tuned to the stimulus dimension [10]. Thus either class of model could be used to alter the image to simulate pattern adaptation effects.

4. An illustration of performance improvements with adaptation

As a proof of concept of how pre-adapting images could facilitate visual performance, we consider the simplest case of adapting to a change in the mean color of an image. Adjustments to the mean color are well predicted by independent multiplicative gain changes in the sensitivity of the three classes of cones, a process known as von Kries adaptation [9]. We examine how these adjustments could affect an observer's ability to detect color targets in the image. In previous studies, we showed that observers are faster at finding an odd color target if they are first adapted to the set of colors defining the background against which the target appears [27]. In the present case, we explore the effects of adapting to backgrounds that differ simply in the average color.

Figure 1 shows the set of colors composing the stimulus in terms of a standard space defined by two chromatic axes that represent relative responses in the long and medium wavelength sensitive cones at constant luminance (L-M) or the responses in the short wavelength sensitive cones at constant luminance (S-(L+M)). These axes thus isolate the component responses of the S cone or the L and M cones for any color and also represent the principal axes along which chromatic information is encoded at early stages of the visual pathway [28]. For our stimulus a background set of colors was displayed on a monitor by showing a dense set of overlapping ellipses with colors drawn from either the S or the L-M axis relative to a mean chromaticity, which in the example shown varied along the +L (reddish mean) or –L (greenish mean) pole of the L-M axis. Target colors were varied to span a range of hues and contrasts varying away from the background axis (unfilled symbols). On a given trial the target was displayed at a random location on the screen, and reaction times were measured for responding whether the target was on the right or left side. The measurements were made after the observer was adapted by viewing random successive samples of the background shown every 250 ms for 3 minutes, and settings were compared after adapting to the actual average color of the background (e.g. +L) or to a background that had a different color than the one on which they had to search for the target (e.g. –L).

Figures 3 plots the reaction times when the observer was searching for a color target that differed from the background, when the background color axis was orthogonal to the axis of the mean color shift. The stimulus colors for this condition are shown in Figure 2. Consider the upper left panel. Here the task was to find target colors that differed from the S axis of the background, and thus to find targets that were defined only by the color difference along the L-M axis. Search times varied from a few seconds for colors close to the background to a few hundred milliseconds for targets that were far removed from the background colors and thus easily "popped out" of the display. These search times are shown by the unfilled symbols for the condition in which the observer was adapted to the average color of the background. The critical comparison is when the observer was instead adapted to the wrong average color. This is shown by the filled triangles and search times are strongly delayed for all targets. The reason for the poorer performance is that by being adapted to the wrong average color the observer's achromatic point was shifted away from the mean

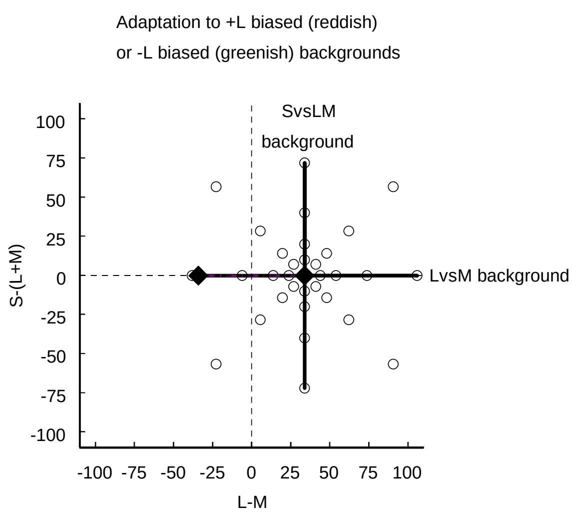

Fig. 1. Set of background (solid line) or target (circles) colors for the visual search task. Observers were adapted to one of two average colors (diamonds), corresponding to the mean of the background or a mean with the opposite color bias.

color of the background so that during the search all colors appeared reddish and consequently less discriminable. That is, the adaptation was optimized for the wrong point in color space. These performance deficits could have been removed by instead adapting the +L image so that its mean coincided with the observer's current state of adaptation (i.e. transforming the image so that average chromaticity was –L). The performance gain from this adjustment is again indicated by the difference in search times going from the unadapted (filled triangles) to the adapted background (unfilled circles).

Figures 4 and 5 show the results for a second set of conditions, in which the adapting axis is now parallel to the axis of the mean color shift. In this case there is no cost or benefit of changing the state of adaptation. Considering again the conditions for the top left panel, the lack of an effect is because the target colors are detected based on the responses along the S-cone axis, yet the adaptation is only changing the sensitivity in the L and M cones. These results illustrate that how adaptively changing images (or observers) will alter performance requires models that are appropriately informed about the actual response changes underlying visual adaptation.

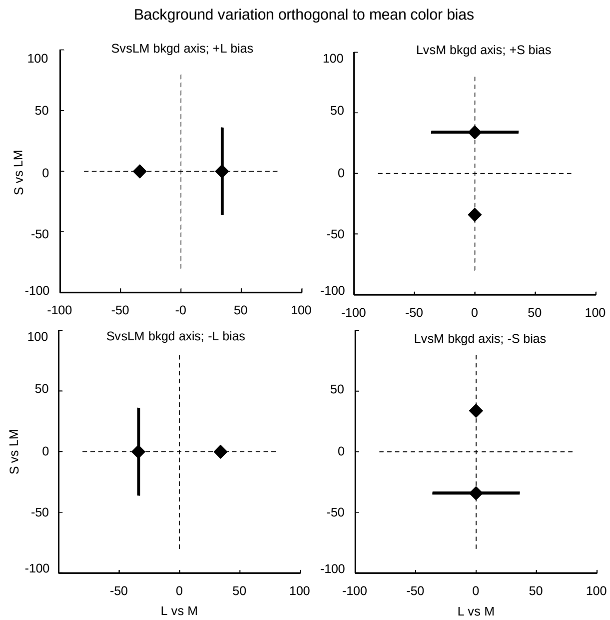

Fig. 2. Mean adapting colors (diamonds) and background colors (lines) when targets had to be detected by the color differences along the axis of the adaptation.

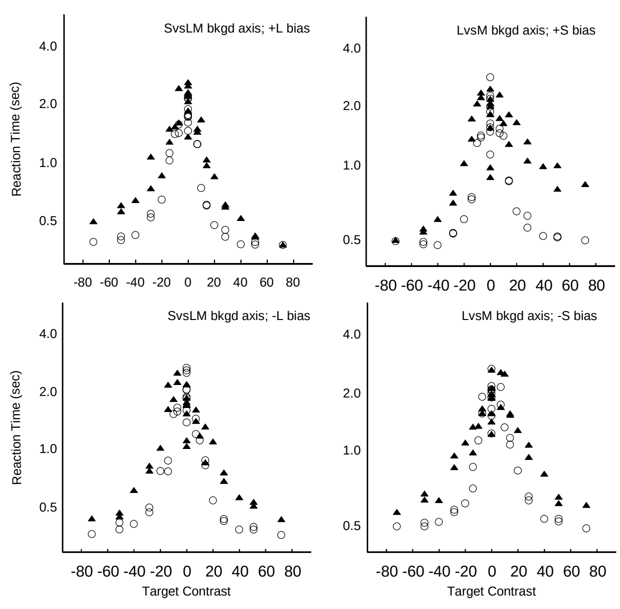

Fig. 3. Search times as a function of the target color contrast relative to the background axis after adapting to the mean background color (circles) or the opposite color bias (triangles).

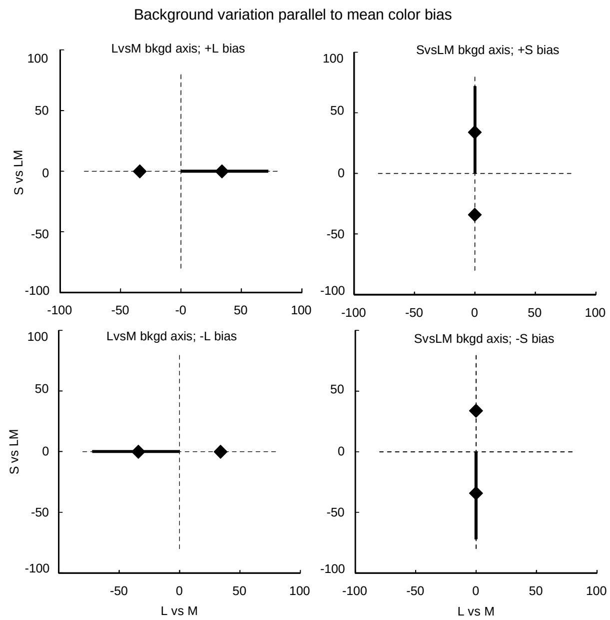

Fig. 4. Mean adapting colors (diamonds) and background colors (lines) when targets had to be detected by the color differences along an axis orthogonal to the axis of adaptation.

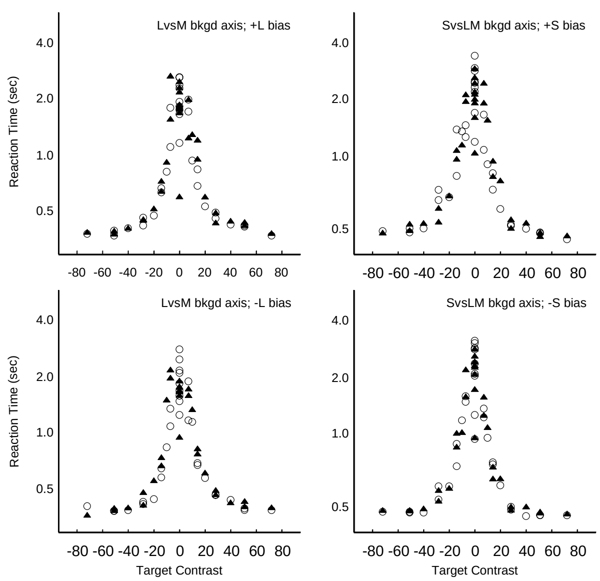

Fig. 5. Search times as a function of the target color contrast relative to the background axis after adapting to the mean background color (circles) or the opposite color bias (triangles).

5 Conclusions

Incorporating the actual coding processes of the human visual system into image processing algorithms can significantly improve image quality (e.g. [29] [30]). In this paper we have argued that visual displays and perceptual performance can potentially be enhanced by mimicking the sensitivity changes that occur whenever the visual system is exposed to an image. If the goals of the observer are known and the stimulus cues that support these goals are well defined, then simulating adaptation may confer little advantage, for in that case the appropriate information can readily be extracted, and many approaches already exist for highlighting task-relevant information in displays. Our approach is instead suited to more open-ended contexts, in which the goals of the observer and the potentials in the image are not always clearly defined. For such contexts it seems likely that the best way to help the observer is to facilitate the very processes that the visual system has evolved to optimize perception.

Examples where this approach might prove especially helpful include tasks that require scanning and interpreting highly biased image sets or highly variable sets. In the former case the observer may be confronted with the equivalent of an "other image effect," in which the bias in the images masks the underlying differences between them, while in the latter the state of adaptation may be unable to track changes in the images at each moment. Another example where pre-adapting the image might prove useful is when the observer is looking for anomalies in image features. As we noted above, one of the postulated functions of adaptation is to enhance the salience of novel properties in the environment, and such effects could be especially relevant in interpreting medical images. In fact, fundus photos of the eye are often pre-adapted to remove the average color bias so that color anomalies are more readily visible, and analogous adjustments could potentially be applied to pre-adapt the spatial structure of images, for example in x-ray images. A further potential clinical application is adapting images to adjust for visual losses owing to injury or disease in order to help the observer to adjust to their changed vision. Finally, simulating realistic adaptation transforms could also prove valuable for extending the sensory limits of vision into different spectral or spatial or temporal frequencies, since these transforms could adjust the stimulus ranges in ways that the visual system is designed to encode.

An obvious question is how quickly an observer can themselves adapt to changes in a display, since this might indicate how beneficial a change in the image could be. While some adaptive adjustments occur very rapidly, others have a remarkably long time course [31]. Thus in principle adapting images may be equivalent to weeks or even months of visual training. The types of perceptual training and visual tasks that could be facilitated are limited now primarily by the limits on current models of perceptual adaptation.

Acknowledgments

Supported by a NASA Nevada EPSCOR Collaborative Pilot Grant and by National Eye Institute Grant EY-10834.

References

1. Webster, M.A. Pattern selective adaptation in color and form perception, in The Visual Neurosciences Volume 2, L.M. Chalupa and J.S. Werner, Editors. 2003, MIT Press: Cambridge. p. 936-947.
2. Webster, M.A. and MacLin, O.H. Figural after-effects in the perception of faces. Psychonomic Bulletin and Review **6**: p. 647-653 (1999).
3. Leopold, D.A., et al. Prototype-referenced shape encoding revealed by high-level aftereffects. Nature Neuroscience **4**: p. 89-94 (2001).
4. Webster, M.A., et al. Adaptation to natural facial categories. Nature **428**: p. 558-561 (2004).
5. Barlow, H.B. Dark and light adaptation: Psychophysics, in Handbook of Sensory Physiology, D. Jameson and L.M. Hurvich, Editors. 1972, Springer-Verlag: New York. p. 1-28.
6. Ohzawa, I., Sclar, G. and Freeman, R.D. Contrast gain control in the cat visual cortex. Nature **298**: p. 266-268 (1982).
7. MacLeod, D.I.A. and von der Twer, T. The pleistochrome: optimal opponent codes for natural colours, in Colour Perception: Mind and the Physical World, R. Mausfeld and D. Heyer, Editors. 2003, Oxford University Press: Oxford. p. 155-184.
8. Laughlin, S.B., Form and function in retinal processing. Trends in Neuroscience **10**: p. 478-483 (1987).
9. Webster, M.A. and Mollon, J.D. Colour constancy influenced by contrast adaptation. Nature **373**: p. 694-698 (1995).
10. Graham, N., Visual Pattern Analyzers. 1989, Oxford: Oxford University Press.
11. Field, D.J. Relations between the statistics of natural images and the response properties of cortical cells. Journal of the Optical Society of America A **4**: p. 2379-2394 (1987).
12. Barlow, H.B. A theory about the functional role and synaptic mechanism of visual aftereffects, in Vision: coding and efficiency, C. Blakemore, Editor. 1990, Cambridge University Press: Cambridge. p. 363-375.
13. Pokorny, J., Shevell, S. and Smith, V.C. Colour appearance and colour constancy, in Vision and Visual Dysfunction 6: The Perception of Colour, P. Gouras, Editor. 1991, Macmilan: London.
14. Werner, J.S. and Schefrin, B.E. Loci of achromatic points throughout the life span. Journal of the Optical Society of America A **10**: p. 1509-1516 (1993).
15. Webster, M.A., Georgeson, M.A. and Webster, S.M. Neural adjustments to image blur. Nature Neuroscience **5**: p. 839-840 (2002).
16. Srinivasan, M.V., Laughlin, S.B. and Dubs, A. Predictive coding: A fresh view of inhibition in the retina. Proceedings of the Royal Society of London B **216**: p. 427-459 (1982).
17. Leopold, D.A. and Bondar, I. Adaptation to complex visual patterns in humans and monkeys, in Fitting the Mind to the World: Adaptation and Aftereffects in High Level Vision, C.W.G. Clifford and G. Rhodes, Editors. 2005, Oxford University Press: Oxford. p. 189-211.
18. Rhodes, G., et al., Fitting the mind to the world: Face adaptation and Attractiveness Aftereffects. Psychological Science **14**: p. 558-566 (2003).
19. Hood, D.C. and Finkelstein, M.A. Sensitivity to light, in Handbook of Perception and Human Performance, Volume 1: Sensory Processes and Perception, K.R. Boff, L. Kaufman, and J.P. Thomas, Editors. 1986, Wiley: New York. p. 5-1 - 5-66.
20. Georgeson, M.A. The effect of spatial adaptation on perceived contrast. Spatial Vision **1**: p. 103-112 (1985).
21. Clifford, C.W.G. Perceptual adaptation: motion parallels orientation. Trends in Cognitive Sciences **6**: p. 136-143 (2002).
22. Walraven, J., et al. The control of visual sensitivity: Receptoral and postreceptoral pro-

cesses, in Visual Perception The Neurophysiological Foundations, L. Spillmann and J.S. Werner, Editors. 1990, Academic: San Diego. p. 53-101.
23. Atick, J.J., Li, Z. and Redlich, A.N. What does post-adaptation color appearance reveal about cortical color representation? Vision Research **33**: p. 123-129 (1993).
24. Barlow, H.B. and Földiák, P. Adaptation and decorrelation in the cortex, in The Computing Neuron, R. Durbin, C. Miall, and G.J. Mitchison, Editors. 1989, Addison-Wesley: Wokingham. p. 54-72.
25. Clifford, C.W.G., Wenderoth, P. and Spehar, B. A functional angle on adaptation in corti cal vision. Proceedings of the Royal Society of London B **267**: p. 1705-1710 (2000).
26. Zaidi, Q. and Shapiro, A.G. Adaptive orthogonalization of opponent-color signals. Biological Cybernetics **69**: p. 415-428 (1993).
27. McDermott. K., et al. Visual Search and eye movements in novel and familiar contexts, In Human Vision and Electronic Imaging XI, B.E. Rogowitz, T.N. Pappas, and S.J. Daly, Editors, SPIE **6057**: p. 60570C-1-12 (2006).
28. MacLeod, D.I.A. and R.M. Boynton. Chromaticity diagram showing cone excitation by stimuli of equal luminance. Journal of the Optical Society America **69**: p. 1183-1186 (1979).
29. Van Hateren, J.H.. Encoding of high dynamic range video with a model of human cones. ACM Trans. Graphics **25**: p. 1380-1399 (2006).
30. Meylan, L., Alleyson, D. and Susstrunk, S. Model of retinal local adaptation for the tone mapping of color filter array images, Journal of the Optical Society of America A **24**: 2807-2816 (2007).
31. Delahunt, P.B., et al. Long-term renormalization of chromatic mechanisms following cataract surgery. Visual Neuroscience **21**: p. 301-307 (2004).