

9

Adaptation and the Phenomenology of Perception

MICHAEL A. WEBSTER, JOHN S. WERNER,
AND DAVID J. FIELD

9.1 Introduction

To what extent do we have shared or unique perceptual experiences? We examine how the answer to this question is constrained by the processes of visual adaptation. Adaptation constantly recalibrates visual coding so that our vision is normalized according to the stimuli that we are currently exposed to. These normalizations occur over very wide ranging time scales, from milliseconds to evolutionary spans. The resulting adjustments dramatically alter the appearance of the world before us, and in particular, alter visual salience by highlighting how the current image deviates from the properties predicted by the current states of adaptation. To the extent that observers are exposed to and thus adapted by a different environment, their vision will be normalized in different ways and their subjective visual experience may differ. These differences will be illustrated by considering how adaptation influences properties which vary across different environments. To the extent that observers are exposed and adapted to common properties in the environment, their vision will be adjusted toward common states, and in this respect they will have a common visual experience. This will be illustrated by considering the effects of adaptation on image properties that are common across environments. In either case, it is the

similarities or differences in the stimuli – and not the intrinsic similarities or differences in the observers – which largely determine the relative states of adaptation. Thus, at least, some aspects of our private internal experience are controlled by external factors that are accessible to objective measurement.

In 2001, a controversial new portrayal of Queen Elizabeth II was unveiled by the painter Lucian Freud. Freud, the grandson of Sigmund, has been hailed as the greatest living portrait artist in England, and had clearly labored carefully over a work that included 70 sittings by the Queen. However, the painting was not well-received. Reviews in the press ranged from muted disappointment (“while she is no longer the heartbreakingly beautiful young woman she was, she is still easy on the eye”) to open hostility (“Freud should be locked in the tower”). Many editorials pointed to distortions in the representation (“the chin has what can only be described as a six o’clock shadow, and the neck would not disgrace a rugby prop forward”). But perhaps, the most telling comment was that “she should have known what to expect,” for the painting bears Freud’s distinctive style and – to the untrained eye – the face depicted seems notably similar to his own self-portrait (Fig. 9.1). Apparently, many in the public saw the painting in a way that the artist did not. In this chapter, we argue that they *literally saw* the painting differently. This is not to suggest that Freud thought that the work actually looked like a faithful copy of the Queen, an error of logic known as the El Greco fallacy (Anstis 1996). Rather, we explore the possibility that Freud might have seen the painting differently, simply because he had spent so much time looking at it.

What might the world look like if we could see it through the eyes of another? Such questions are central to the debate over the nature of perceptual



FIGURE 9.1 A portrait of *Her Majesty Queen Elizabeth II* (2001) by the artist Lucian Freud, and his self-portrait *Reflections*.

experience and sensory qualia, or what it “feels” like to see. Because we have access only to our own private experience, we cannot directly observe whether it is similar in others. A classic example of this limitation is Locke’s inverted spectrum (Locke 1689, 1975). Even if two observers completely agree on how they label the hues of the spectrum, we cannot be certain that their experiences agree, for the subjective sensation of redness in one might correspond to the sensation of greenness in the other. Arguments about phenomenology must instead rely on inferences from indirect observations. For example, arguments against a phenomenally inverted spectrum have pointed out that this possibility would be inconsistent with asymmetries in the properties of color perception (Hardin 1997; Palmer 1999).

In this review, we consider how the nature of subjective experience is constrained by the processes of sensory adaptation. Adaptation adjusts visual sensitivity according to the set of stimuli an observer is exposed to. As the many chapters in this book illustrate, such adjustments are a built-in feature of visual coding and probably regulate most if not all aspects of visual perception. Indeed, adaptation may represent a fundamental “law” of cognition and behavior, a point most forcefully argued by Helson (1964). Here we focus on how specific presumed properties of visual adaptation might be expected to influence visual phenomenology. Studies of adaptation aftereffects have shown that changes in the state of adaptation have dramatic consequences for how the world looks like. The states of adaptation may therefore play a fundamental role in determining whether the world looks the same or different to others.

9.2 Adaptation and Response Normalization

The use of information theory has provided major insights into our understanding of sensory coding. By understanding the statistics of the environment and relating those statistics to response properties of sensory neurons, we have learned that these sensory neurons are providing a highly efficient representation of the environment (Atick 1992; Field 1994; Simoncelli & Olshausen 2001). It seems reasonable to assume that the processes of perceptual adaptation contribute to this efficiency in coding (Wainwright 1999). To understand how the phenomenology of adaptation might bear on such coding, we first consider the influence of adaptation on individual neurons and then on the distribution of responses across neurons.

Neurons have a limited dynamic range, and because they are noisy, can reliably signal only a relatively small number of response levels (Barlow & Levick 1976). To maximize the information a neuron can carry, these levels should be matched to the distribution of levels in the stimulus. This principle closely predicts evolutionary adaptations such as the sigmoidal shape of a neuron’s response function (Laughlin 1987). Most points in a scene have a brightness and color that are close to the modal level, and thus the optimal response function should be steep near the mode, to allow fine discrimination among frequently

244 FITTING THE MIND TO THE WORLD




FIGURE 9.2 Fitting the response to the stimulus. (Left panel) For a Gaussian distribution of stimulus levels, the optimal response curve will be steepest near the mean, where stimuli are frequent, while asymptoting at the tails, where signals are rare. (Right panel) For an opponent mechanism the response curve signals deviations (positive or negative) from the mean; this opponency can also be realized in all-positive response curves by splitting the response into separate on and off mechanisms.

occurring stimulus values, while shallow at the tails, where signals are rare. This effectively expands the representation of data near the modal level and compresses those data near the outliers (Fig. 9.2).

The same considerations also predict the need for short-term adaptations, since the variation within any scene will be less than the variation between scenes. Therefore, a system that can “float” the sensitivity range can maximize the information carrying capacity of a neuron. An obvious example is the enormous variation in the average light level during the course of the day. The intensity variations within a scene (in the range of 300 to 1) are much lesser than the intensity variations across scenes (on the order of 10^{14} to 1). Therefore, a system that can recalibrate to the individual scene can reduce the required dynamic range by several orders of magnitude. Without this adaptation, any given neuron with its limited dynamic range, would be silent or saturated most of the time (Walraven *et al.* 1990).

By adjusting to the average stimulus, the visual system could represent information in the form of deviations from the average. This gives special importance to the mean because it defines the reference point to which other responses are now relative. One way to realize this reference in the neural response is to use an opponent code, in which the responses can be of opposite sign. For example, the intensity response of a mechanism can be recoded so that there is zero response to the mean intensity, and darker or brighter stimuli are represented by negative or positive values, respectively (Fig. 9.2). Opponent processing is a hallmark of color vision: color-opponent mechanisms receive inputs of opposite sign from different cone types and thus their outputs represent a comparison of the relative activity across the cones (De Valois 2003). It may be that opponent processing is more generally a central property of perception because of the general need to make comparisons (Hurvich & Jameson 1974). A consequence of opponency is that the neuron is silent to the average. Thus, a “red versus green”

Q: Please confirm if a small addendum to this sentence is allowed so as to read:

"This split code can improve efficiency by increasing the signal-to-noise ratio over a pair of neurons that allow both the neurons to span the full dynamic range so that each neuron codes for only one half of the input range, while opponency itself increases the metabolic efficiency by greatly reducing the average firing rate of cells."

mechanism does not respond to "white," or importantly, to the average color that it is exposed to. This average is thus represented only implicitly, by the absence of a signal. Note that within a single neuron, responses of opposite sign are relative to the neuron's background activity. However, these opposing responses may instead be instantiated within separate "on" and "off" mechanisms. **This split code can improve efficiency by increasing the signal-to-noise ratio over a pair of neurons that instead both spanned the full dynamic range,** (MacLeod & von der Twer 2003) while opponency itself increases the metabolic efficiency by greatly reducing the average firing rate of cells.

To realize its full capacity, a neuron's operating curve should be matched not only to the average stimulus, but also to the range of stimulus levels or available contrasts. A clear example of this optimization is in color coding. Color vision depends on comparing the responses across different classes of cone. Yet, because the spectral sensitivities of the cones overlap, this difference signal is necessarily smaller than the range of available luminance signals (which instead depend on adding the cone signals). If the post-receptoral neurons encoding luminance and color had similar dynamic ranges, then they would again be silent or saturated most of the times. Instead, chromatic sensitivity is much higher than luminance sensitivity, consistent with matching responses to the available gamut (Chaparro *et al.* 1993). However, in this case again, the environment can vary in the range of stimulus contrasts, and thus short-term adaptations would again be necessary if the neurons are to be appropriately tuned to the scenes before us. This form of adjustment, known as contrast adaptation, is well-established psychophysically as well as in individual neurons (Webster 2003). Thus, for example, sensitivity to contrast is reduced in the presence of high contrast stimuli (though the precise form of the response changes or their functional consequences is less clearer than for light adaptation). Our first point then, is that it is plausible to expect adaptation to play a pervasive role in normalizing neural responses and that these adjustments should operate in similar ways across observers. Whether or not two observers have similar subjective experiences should therefore be, in part predictable from whether or not they are under similar states of adaptation.

At least in the early stages of the visual system, it is common to assume that information is encoded by channels that are selective for different values along a stimulus dimension. For example, color is initially encoded by three channels – the cones – that differ in their selectivity to wavelength. How should the responses across these channels be distributed? Assuming that neurons involved in different computations have a roughly similar dynamic range, we can predict that if we shift to a new set of axes (e.g. an opponent system), then we can expect the magnitudes along the different axes to be roughly similar (Fig. 9.3). This predicts that the channel gain should be inversely proportional to the strength of the stimulus component for which they are selective. One example of this principle in color vision is the gain of signals derived from S cones. The S cones make up only a small fraction of the total number of cone receptors and the wavelengths to which they respond are more strongly filtered by the lens

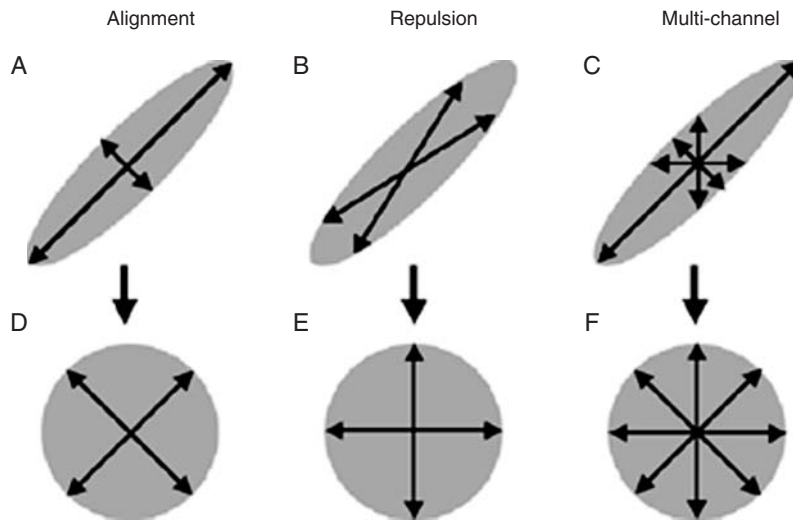


FIGURE 9.3 Three ways in which a set of channels might selectively adapt to sphere the responses. (A) and (D) alignment of the response axes along the principal stimulus axes and gain control; (B) and (E) repulsion of the axes relative to the principal axis; (C) and (F) independent gain changes within multiple channels tuned to different stimulus directions.

and macular screening pigments, yet their signals are greatly amplified in the visual cortex so that the response to different hues is more effectively “spherized” (De Valois *et al.* 2000).

A second example of matching channel gain to the stimulus distribution is provided by the spatial statistics of images, which have less contrast at fine scales (high frequencies) than coarse scales (low spatial frequencies). Cortical mechanisms tuned to different spatial scales may vary in sensitivity in such a way that compensates for the stimulus bias so that the response across scale is the same on an average (Field & Brady 1997). Both of these examples could reflect on the evolutionary adaptations of the visual system in response to more or less stable attributes of the visual environment. Yet in both cases, we now consider how these stimulus properties can routinely change because of changes in the environment or the observer, and thus short-term adaptive adjustments would again be important for maintaining the balance across the channels. In fact, the processes that adjust each neuron to the average stimulus level it is exposed to, will serve to balance the responses across neurons. Thus, our second point is that adaptation will normalize visual responses to adjust to the specific biases in the observer’s environment.

A final prediction is that the responses of different neurons should be as independent as possible. If two neurons are redundantly carrying the same information, then together they will require greater information capacity than if the data were represented independently. For example, the signals from different

cone classes are highly correlated. Postreceptoral neurons remove much of this redundancy by recoding the cone signals into their sums or differences (Buchsbau & Gottschalk 1983). In a similar way, the center-surround organization of retinal receptive fields can be viewed as a strategy for removing the correlations between the receptor responses to nearby regions of space, which tend to have similar brightness and color (Srinivasan *et al.* 1982). However, although the channels may be largely uncorrelated when considering the population of all images, any given image or environment may have relatively strong correlations. Furthermore, there are likely to be correlations between different stimulus dimensions (e.g. between brightness and color) that will vary between environments. Thus, a channel structure that will allow the system to dynamically tune its responses to different environments will provide a means of making maximal use out of the limited dynamic range of the system.

To remove the correlations between a set of channels, it is obviously not enough to equalize their responses – instead, the actual tuning of the channels must change, and there is evidence that adaptation can influence not just the gain, but also the stimulus selectivity of individual neurons (Carandini *et al.* 1997; Movshon & Lennie 1979; Dragoi *et al.* 2002). How might adaptation adjust the neural representation to achieve independence? It turns out that there are many possible strategies, and these are illustrated in Fig. 9.3 for a pair of mechanisms encoding an elliptical stimulus distribution. One way would be to rotate the preferred axes of the mechanisms to align them with the principal axes of the stimulus variance (Fig. 9.3(A)). This process, called Principal Components Analysis (PCA), has been used to predict the realignment of the cone axes along separate luminance and color-opponent channels (Buchsbau & Gottschalk 1983), and a rotation followed by scaling (Fig. 9.3(D)) would effectively sphere the responses. However, note that after sphering, the vectors will remain independent for any orthogonal rotation. Thus, decorrelation alone places weak constraints only on the choice of axes. An alternative principle is to select the channels in order to capture the separate components contributing to the distribution. This method, known as Sparse Coding or Independent Components Analysis, is more effective for non-elliptical stimulus distributions (e.g. composed of multiple ellipses from multiple sources), and has been found to closely predict the receptive field properties of cells in striate cortex (Olshausen & Field 1996). Both of these processes may be important in longer-term evolutionary adaptations of visual coding, and it is an intriguing question whether similar types of adjustments – and specifically, adjustments that align the channel axes along the stimulus axes – can occur through short-term adaptation. The signature of these might be changes in the “labels” carried by the channels. However, short-term visual aftereffects have instead been interpreted in terms of response changes within channels whose labels are fixed.

Barlow and Földiák (Barlow & Földiák 1989; Barlow 1990) suggested that perceptual adaptation might reflect on mutual inhibition that builds up between channels whenever they respond together, thus biasing their preferred axes by setting up a mutual repulsion between them (Fig. 9.3(B and E)). Note that this

results in a different pair of channels than the pair predicted by PCA, because the response vectors are biased relative to the stimulus axes rather than aligned with them. Their model provides a physiologically plausible mechanism for decorrelation and could in particular account for the phenomenon of “contingent adaptation.” For example, in the McCollough Effect (McCollough 1965), color aftereffects are elicited that are contingent on the orientation of the adapting grating, and the sensitivity changes may reflect upon adaptive adjustments that are designed to remove the correlation between color and orientation.

However, a problem arises when the set of data axes and the number of channels (i.e. vectors) is overcomplete, or in other words, when a larger number of output vectors is used to span a smaller number of input dimensions (e.g. Fig. 9.3(C)). For example, for the plane in Fig. 9.3, it is not possible to have an overcomplete set of vectors in addition to keeping them uncorrelated. Field (in preparation) argues that with certain forms of non-linearity the family of vectors can be both uncorrelated and overcomplete. The details are beyond the scope of this chapter. However, the issue of overcomplete coding and its relationship to optimal mapping is an important issue. The number of neurons in the visual cortex is at least an order of magnitude larger than the number of inputs into it. For instance, signals from the three cone classes appear to be carried within more than three types of color channels in the cortex (Lennie 1999), and thus simple linear operations cannot remove the correlations between them. Moreover, in practice, the aftereffects predicted by linear decorrelation are similar to the changes predicted by an independent adaptation within multiple channels (Fig. 9.3(C and F)). That is, both predict a selective loss in sensitivity to the principal axis of the stimulus even though the basis for this loss is entirely different (Webster & Mollon 1994). Thus it has proven difficult to distinguish between different models of the adaptation. Nevertheless, in each of these cases the selective aftereffects can be thought of as spherizing the representation of the environment, and thus the principle of spherizing remains a powerful predictor of the consequences of adaptation for visual experience.

9.3 Adaptation to a Colored Environment

The implications of these adaptive adjustments for the problem of other minds can be illustrated by considering the simple case of adapting to the average color of a scene. Suppose that two otherwise identical observers are placed in a pair of identical rooms, except that in one room the average color is white while in the other it is red (Fig. 9.4). Adaptation will adjust the gains of the cones so that their average responses are equated in each room, a process that in the case of color vision is known as von Kries scaling (von Kries 1970). As a result, over time both observers should report that their room is more nearly achromatic, and if the adaptation were complete then the color difference would be completely factored out. However, if they were suddenly able to look through

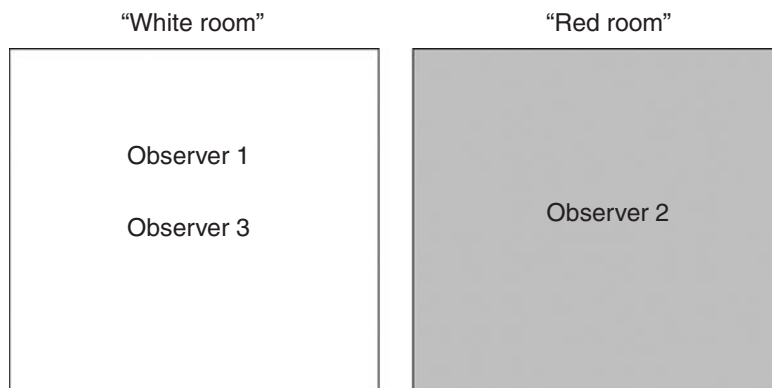


FIGURE 9.4 Adaptation to different colored environments. Observers 1 and 2 are assumed to be physiologically identical but will experience color differently because they are adapted to “white” and “red” worlds. Observers 1 and 3 are assumed to have different spectral sensitivities but will have similar color experiences because they are adapted to the same world.

each others’ eyes then their perceptions would not agree, for what the observer in the reddish room perceived as white, would appear red to the other. Thus, two observers who are intrinsically the same will perceive the same stimulus differently if they are adapted to different environments.

Now suppose a third observer is placed in the white room, and that the spectral sensitivity of this observer differs from that of the first. For example, this individual might have a higher density of lens pigment, reducing the amount of shortwave light reaching their retina. However, von Kries scaling should again adjust their perception so that the room appears achromatic. The first and the third observer would therefore agree about the color of the room, and in this sense would have a common perceptual experience. Thus, even if two observers are physiologically different, their phenomenal experience may be similar in important ways if they are adapted to the same environment.

These two cases illustrate a fundamental asymmetry in the adaptation. Specifically, the visual system is adjusting to the properties of the environment, or as stated by Hecht, “the stimulus adapts the organism to itself” (Hecht 1923, p. 577). This means that in order to predict whether the observers have common or different perceptual experiences it may often be more important to characterize the environment than the observer.

It turns out that we do not have to imagine these examples in the minds of others, for we can each readily experience them in our own visual systems. The afterimages we experience when we look at a dark or colored form, and then shift our gaze to a uniform field are the lingering aftereffects of an adjustment that serves to normalize sensitivity to the brightness and color of the adapting stimulus. As the stimulus changes, our color perception is shifted accordingly. Alternatively, as we look at a uniform field we do so through a retina which is

250 FITTING THE MIND TO THE WORLD

highly inhomogeneous, yet our subjective impression often gives little hint of this pronounced physiological variation. It may be that this is in part because the physiologically different areas of the retina have nevertheless been normalized through adaptation to roughly the same external stimulus.

9.4 Varieties of Perceptual Constancy

It is instructive to reconsider color experience within the two rooms of Fig. 9.4 from the perspective of perceptual constancy. Constancy refers to the stable perception of object properties despite changes in the viewing conditions. Most discussions of constancy focus on how the visual system can discount or compensate for variations in the stimulus. Yet the notion of constancy can also be applied to the problem of discounting variations within the observer, and by extension, can be generalized to variations between observers. These different aspects of constancy are illustrated in Table 9.1 and are discussed in turn.

9.4.1 *Compensating for Variations in the Stimulus*

Many forms of stimulus constancy, such as veridical judgments of size despite changing distance or of shape despite changing viewpoint, may not involve adaptation in a direct or obvious way. However, color vision is a case in which adaptation plays a widely recognized role. Under most natural viewing conditions the perception of color is more closely correlated with an object's reflectance properties than with the distribution of wavelengths the object is currently reflecting. The latter is a product of both the reflectance function and the incident illuminant, and the problem for color constancy is how to factor out the contribution to color from the illuminant (Lennie & D'Zmura 1988; Zaidi chapter, this volume). The human visual system is only approximately color-constant and employs many different strategies to achieve this, but a simple mechanism that works to promote color constancy is von Kries scaling of the cone signals (Brainard & Wandell 1992). In Fig. 9.4, suppose we change the white room into a red room simply by changing the illuminant. We have already seen that an observer in this room will renormalize to the new mean

Table 9.1 Forms of perceptual constancy. Constancy can compensate for changes in the stimulus (e.g. a change in the color of the illuminant) or the observer (e.g. a change in spectral sensitivity with aging), and thus could also compensate for environmental or physiological differences between observers

	Environmental variation	Observer variation
Within-observer	E.g. illuminant change	E.g. aging lens
Between-observers	E.g. two observers under two different illuminants	E.g. two observers that differ in lens density

chromaticity, so that the average color appears achromatic. As a result of this adjustment, the spectral bias introduced by the new illuminant is partially discounted. For example, for the surfaces with a flat reflectance function, the spectrum of the reflected light will mirror the illuminant, and thus without adaptation, would appear as the color of the illuminant. Yet, to the extent that adaptation renormalizes the cone signals to “flatten” the response to this spectrum, the surface itself will appear spectrally flat (“gray”) and thus, will be the same under either illuminant.

In what sense does observer 1 maintain constant color experience if the lighting is changed to red, when we had argued previously that (as observer 2) their color perception does change when the room itself is red? The difference is because we have assumed that their subjective experience refers to the object reflectances. In either case, von Kries adaptation will adjust to the average chromaticity. Yet when the lighting changes, this adaptation removes the illuminant color and thus preserves the same response to the same reflectance, while in the previous case it is the objects themselves that have changed, and thus the perception of object reflectance changes. [Note that to illustrate these points we have considered only a simple adjustment to the mean chromaticity. In actual practice, the visual system can exploit differences in how color signals vary with changes in reflectance versus lighting in order to disambiguate between them, and thus can exhibit better constancy for both reflectances and illuminants than that predicted by von Kries adaptation alone (Golz & MacLeod 2002)].

In general, adaptation should lead toward constancy when it factors out extraneous sources of variation in the stimulus. Again, in the case of color or lightness constancy, these arise because of variations in the spectral content or intensity of the illuminant. The visual system may also exhibit contrast constancy, and perhaps blur constancy, in scenes with reduced visibility as when we view them through fog (Brown & MacLeod 1997). All of these cases represent situations in which the extraneous variable is the light medium. Are there analogous cases in other domains of perception, in which adaptation can compensate for irrelevant sources of stimulus variance? For example, in face or speech perception, the stimulus is shaped by both individual factors (e.g. identity, fitness, or expression) and more global factors (e.g. familial or ethnic characteristics). If adaptation allows the visual system to parse out and compensate for these more global population characteristics, then it is meaningful to suppose that it provides a form of “face constancy” for individual characteristics. For example, discounting the average characteristics of a population (by normalizing for these characteristics) may allow the observer to maintain stable perceptual criteria (e.g. “average” features) for judging properties of individual faces such as attractiveness (Rhodes *et al.* 2003).

9.4.2 *Compensating for Variations in the Observer*

What will happen if observer 1 changes over time to become observer 3, because properties of their visual system change? Here again, we see that adaptation

should adjust to compensate for these physiological differences. From the perspective of constancy, this means that the observer should continue to maintain a stable perception of the world even though they have themselves changed. This is accomplished by factoring out extraneous sources of variation which are now due to the observer. This recalibration may represent one of the most important and general forms of perceptual constancy. The genetic specification of visual pathways is unlikely to provide more than a broad outline for coding, leaving much to be shaped by experience. The problem of calibrating the channel structure is obviated by the processes of adaptation, since these adjust visual sensitivity in order to match it to the environment. Thus, the adaptation represents a form of “error-correction” in the observer (Andrews 1964). In the same way, these processes should correct for the errors that could arise in the course of normal development and aging, during which there are profound changes throughout the visual system (Werner 1998). As we describe below, studies of color perception across the lifespan provide powerful evidence for perceptual constancy despite the dramatic changes in visual pathways with aging.

9.4.3. *Interobserver Constancy*

If we accept that adaptation can maintain a stable perceptual experience despite changes within an observer, then it is perhaps not too far a leap to suppose that the same processes can contribute to constancy between observers. This “inter-observer constancy” means that the world should look similar to different observers in important ways. The states of adaptation may not speak to the nature of qualia. For example, they do not require that what white feels like to one will be similar to another. However, they do strongly influence what stimulus appears white, or more generally, what stimuli appear perceptually neutral to an observer, and thus whether these perceptual landmarks are equivalent in others. These shared perceptions are possible when observers are adapted to similar worlds, or when the adaptation factors out the differences between their worlds (e.g. differences in illumination); but should break down when observers are instead normalized for different properties of their (different) environments.

9.5 Variation and Constancy in the Visual Environment

9.5.1 *Color and the Natural Environment*

In the previous sections, we explored the consequences of adaptation by asking how we might be adapted if we lived in different colored rooms. But to what extent do our color worlds actually vary? A number of studies have examined the color statistics of natural images (Burton & Moorhead 1987; Webster & Mollon 1997; Párraga *et al.* 1998; Ruderman *et al.* 1998). Short-term color differences routinely arise from differences in illumination, but more long-term differences also result from differences in the settings themselves. One principal source of variation is between arid and lush scenes, which can contribute to

differences both across environments and within the same environment over time. For example, Fig. 9.5 shows measurements from a set of scenes from a single valley in rural India during the monsoon (wet) and winter (dry) seasons (Mizokami *et al.* 2003). The changes in vegetation shifted the mean color from green to yellow, and shifted the principal axes of the distributions from near the S axis of cone-opponent space to a bluish-yellowish axis intermediate to the S and LM axes. Note that these mean color shifts were large compared to the differences in the illuminants during the two seasons. Thus, different environments can in fact vary widely in their color properties.

Are these differences large enough to induce different states of adaptation? Webster and Mollon (1997) exposed observers to a sequence of colors drawn from different measured scenes, and then used a matching task to measure changes in the appearance of a set of test colors. Adaptation altered the perceived color by adjusting to both the average color in the scene and to the contrast axis defining the scene. For example, after viewing the colors from an arid yellowish scene that varied primarily along a blue-yellow axis, the mean color

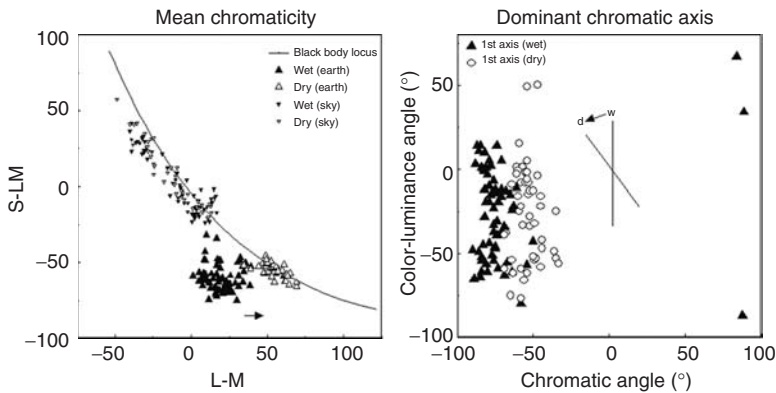


FIGURE 9.5 Seasonal variations in color distributions. Each point describes the color properties of one image taken within the same vicinity during wet or dry seasons in Nashik district, India. The color values are represented according to a scaled version of the MacLeod-Boynton/Derrington-Krauskopf-Lennie color space, which forms a sphere with “gray” at the origin. Color coordinates correspond to the distance from the origin along the three cardinal axes of postreceptoral color coding (representing variations in luminance or in color along axes that vary in activity in L versus M cones or the S cones versus the sum of L and M cones). (Left) Average color in the scenes. Seasonal changes in vegetation cause the average color to shift along the L-M axis from green to yellow; these shifts in reflectance are large relative to the seasonal differences in illumination (shown by the small symbols for sky); (Right) Principal axis of the image colors. The principal directions of the color distributions also rotate over time, varying primarily along a bluish-greenish axis (the vertical S versus LM axis in the color space) during the wet monsoon, while along a bluish-yellowish axis (~ -60 deg) during the dry winter.

for this distribution appeared whiter, and there was a selective loss in the perceived saturation of blues and yellows compared to other hues (Fig. 9.6). This is consistent with an adjustment to the mean color through chromatic adaptation and with an adjustment to the relative variance in different color directions through contrast adaptation. These selective shifts in color appearance were also found with a hue-scaling task for colors presented on spatially varying backgrounds (Webster *et al.* 2002*b*). Thus, the color differences across scenes are in principle sufficient to hold observers under different adaptation states. However, it ultimately remains to be demonstrated whether individuals who actually live in different settings would show differences in color perception that can be linked to different states of adaptation (Webster *et al.* 2002*c*). The strongest test of this possibility might be to measure the white point for populations immersed in characteristically different environments.

9.5.2 Adaptation to Faces

A second example where stimuli might vary across environments and thus shape perception in different ways is in the perception of faces. Adaptation to




FIGURE 9.6 Adaptation to natural color distributions. (Left) color distribution measured from an arid scene in the Sierra Nevada mountains, plotted within the scaled MacLeod-Boynton diagram. Points show the chromaticities of points sampled from different locations within the same scene. The set of colors in the scene is largely confined to a bluish-yellowish (sky and earth) axis varying around the yellowish mean, lying at an orientation of roughly -55 deg within the color plane. (Right) shifts in color appearance after adapting to a rapid sequence of colors drawn from this distribution. Test colors (unfilled circles) were a set of stimuli centered on the adapting mean, and all appeared yellowish under neutral adaptation (to the reference white). After adaptation (filled circles), the average color instead appeared white (consistent with normalizing for the adapting mean), and stimuli along the blue-yellow axis appeared less saturated (consistent with normalizing for the adapting contrasts). Open triangles are the color changes that would be predicted if adaptation adjusted only to the average color, and not the color contrast, through independent gain changes within each cone class (von Kries scaling).

faces is reviewed in detail in other chapters of this volume. Here we focus only on the evidence that these adaptation effects are large for the natural variations that characterize human faces and that they can strongly affect the kinds of perceptual judgments that we normally make when we look at faces. These include such categorical judgments as the identity, gender, expression, or ethnicity of an individual. To explore these, Webster *et al.* (2004) took images of actual faces from different categories (e.g. a male and female face), and then morphed between them to form a finely graded series of intermediate images. Subjects varied these images to estimate the boundary between the two perceived categories (e.g. the point at which the image stopped looking “male” and started looking “female”). These ratings were strongly and rapidly biased by prior adaptation to one of the original faces. Thus after viewing the male image, intermediate morph levels were more likely to be perceived as a female, and consequently the category boundary shifted toward the adapting image (Fig. 9.7). These perceptual shifts are consistent with a renormalization of face perception so that the face we are currently exposed to appears more neutral.

Similar aftereffects were found for faces that differed in ethnicity or expression. Moreover, the aftereffects also transferred across different faces. Thus exposure to a sequence of female faces biases the gender boundary between male and female images that were not part of the adapting set. Such results suggest that adaptation can – and probably routinely does – strongly shape our face perception. Again, faces are interesting in this regard precisely because they vary across individuals. Yet these variations are not random. We are each exposed to a different diet of faces because we live in environments peopled by different distributions of individuals, ages, genders, and ethnic groups. To the extent that our perception is normalized for the specific characteristics of these distributions, our perception of faces should differ.

9.5.3 Color Vision and Aging

The preceding two examples illustrate the possibility that individuals see the world differently because they are adapted to different worlds. Can we identify reasonable instances where physiologically different individuals may perceive in similar ways because they are adapted to the same world? One instance can again be drawn from color vision, and specifically, from measures of color appearance across the life span.

Sixty years after Wright (1928–29) discovered that age-related brunescence (i.e. yellowing) of the lens of the eye causes rather substantial changes in color matching, he posed the question: “Why do the colors of familiar objects look exactly the same to me now as they did when I was a boy?” (Wright 1988, p. 138). In addition to the lenticular pigment changes that Wright knew about, there are numerous other age-related changes in the eye, the photoreceptors and the visual pathways. Of most interest, however, is the lens itself, for its effect on the visual stimulus is well-understood; it selectively attenuates short wavelengths, and as a consequence, the relative excitation of the three different classes of cone photoreceptors. The effect

256 FITTING THE MIND TO THE WORLD

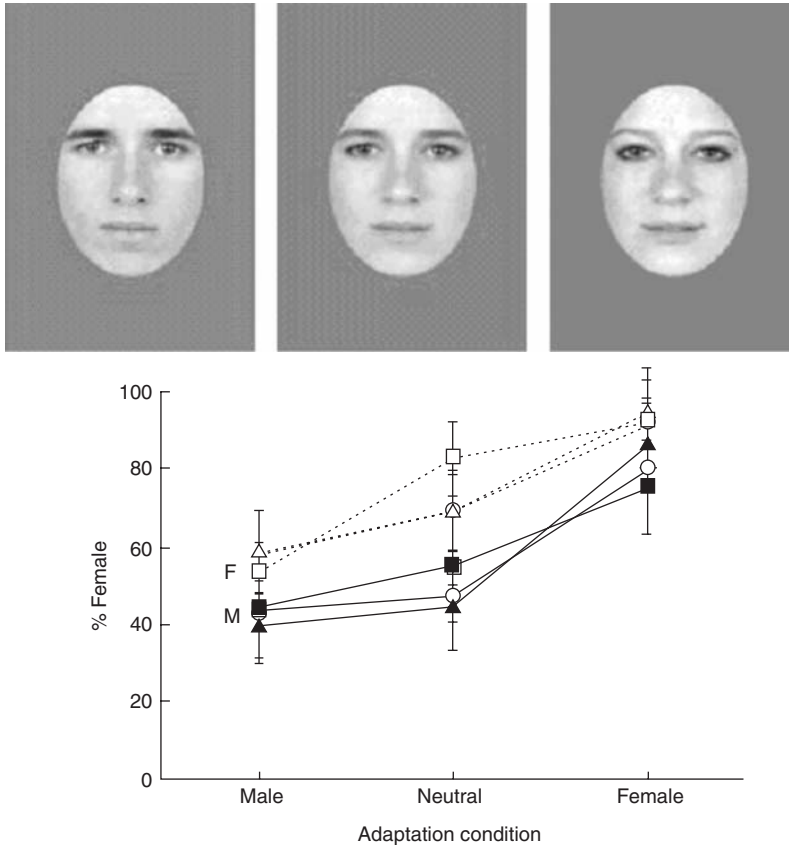


FIGURE 9.7 Adaptation to natural variations in faces. Subjects were shown a graded set of 100 images formed by morphing between a male and female face, and chose the point along the continuum that divided the images into “male” or “female”. The three images at the top show an example of one of the male and female face pairs and the intermediate morphed image between them. Curves plot the settings for a female (F, open symbols) or male (M, closed symbols) subject, each tested with three different morph pairs (square, circle, triangle). Note that each subject chose a boundary closer to their own gender. After adapting to a male (female) face, the blended images appear more female (male), shifting the perceived neutral point toward the adapting gender.

is quite substantial, as shown in Fig. 9.8 where the change in excitation of S-, M-, and L-cones is plotted as a function of changes associated with age (due to increasing lens density) and phases of daylight illumination (or correlated color temperature). While the changes in cone excitation are similar for aging and changes in the color of the ambient illumination, they take place on different time scales, and the visual system adapts to them on different time scales, as discussed in the next section.

Wright’s (1988) suggestion that colors appear the same over a wide range of ages, despite changes in the retinal stimulus and the visual system, was based on introspection, but has been supported by a variety of color appearance studies.




FIGURE 9.8 Relative excitation of the S-, M-, and L-cones is plotted as a function of two variables, age (or the changes in the optical density of the lens correlated with age) and phases of daylight illumination (correlated color temperature).

Consider, for example, the locus in color space of the stimulus that appears achromatic. This stimulus represents the balance point of all color mechanisms and so, ought to be sensitive to changes in that balance in any direction in color space whether those imbalances originate in receptor or postreceptor pathways. In Fig. 9.9, arrows show predictions which result from an experiment in which subjects were asked to adjust a light mixture until it appeared achromatic. The origin of the arrow denotes the location of a theoretical achromatic point for a 10-year-old observer, the ratio of short-wave and long-wave lights adjusted to




FIGURE 9.9 The origin of the arrow represents the location of a theoretical achromatic point for a 10-year-old observer plotted in the CIE 1976 U', V' chromaticity diagram (a standard color space in which equal distances are roughly equal perceptual steps). The arrowhead shows the expected shift in this point (an additive mixture of short- and long-wave monochromatic light) resulting from changes in the ocular media from age 10 to 80. The overlapping data points represent average achromatic loci of different age groups. (Data from Werner and Scheffrin 1993.)

appear devoid of all hue. Because aging of the lens progressively reduces the amount of short-wave light that reaches the retina, more of it will be needed as the person becomes older. The head of the arrow shows how the stimulus would have to be changed to maintain the same retinal stimulus for the same observer at age 80. In other words, the achromatic point at age 80 would be quite blue at age 10. This is not what happens. As shown by the data points representing the average of about seven observers in each age group, the stimulus that appears achromatic is relatively constant. Stability of the achromatic point would not be possible in a static visual system, but is obtained by the visual system renormalizing itself for changes in the retinal stimulus and any age-related physiological changes.

Lens density increases across the life span, but the rate depends on many factors. Especially important in modulating lens aging is the exposure to ultraviolet-B (UV-B) radiation. Higher exposure to UV-B radiation due to more sunlight exposure or living at latitudes closer to the equator is associated with more rapid lens aging (Young 1991). This implies that the relative reduction in S-cone stimulation associated with normal aging is accelerated in people residing at lower latitudes. The color lexicon is also different in many cultures closer to the equator in that some lack separate terms for “green” and “blue.” The combined term is sometimes called “grue.” Lindsey and Brown (2002) have hypothesized that this is due to their increased exposure to UV-B radiation leading to premature lens brunescence and a loss in the ability to see “blue” due to a reduction in short-wave light reaching the retina. They demonstrated that when the short-wave energy in stimuli was reduced in a manner that simulated an older lens, individuals were significantly less likely to use the color term “blue.” This result suggests that adaptation to age-related changes in the lens is not complete even within the normal range before cataract formation. However, we now know that long-term adaptation to brunescence of the lens may be different from short-term adaptation in simulated aging. To test this hypothesis, Hardy *et al.* (2005) examined color-naming patterns in native English speakers with a wide range of ocular media densities (a logarithmic scale); ten were younger (18–35 y.o.) and ten were older (65–85 y.o.), but all had normal color vision. Lens density was measured for each individual prior to tests in which 40 stimuli (simulated Munsell chips) were presented on a computer screen as 2° disks on a gray (illuminant C) background. The same set of stimuli was also presented but with a modification such that the older subjects viewed stimuli that simulated chips filtered through the ocular media of an average younger observer, while younger subjects viewed stimuli that simulated chips filtered through the ocular media of an average older observer. As expected from Lindsey and Brown (2002), younger subjects used “blue” significantly less often in the simulated aging condition. However, for the first set of stimuli, older subjects used “blue” just as often as younger subjects did. In fact, color naming in the younger and older groups was virtually identical for the first set of stimuli. As shown in Fig. 9.10, there was no correlation between the proportion of “blue” responses and density of the ocular media. Thus, while lens brunescence is likely to be accelerated in individuals with higher UV-B exposure, this does not explain the existence

AQ: Please clarify the year of publication.




FIGURE 9.10 The proportion of blue responses to 40 simulated Munsell chips presented on a computer screen is plotted as a function of ocular media density. The density range shown here spans the lowest values typically found in an adult population to values typically diagnosed as early-stage cataract.

of “grue” in high UV-B areas. Indeed, the visual system adapts to these changes in the retinal stimulus, at least over a large range, to maintain constancy of color appearance across the life span.

9.5.4 Adaptation and Blur

A second example where adaptation to a common environment may factor out large intrinsic differences between observers is in the perception of image blur. Blur is an important dimension of image quality, and like face recognition, is a stimulus that we make intuitive and natural judgments about all of the time. There is no question that blur is also a property to which the visual system is constantly adjusting. However, most studies examining these adjustments have focused on the accommodative changes in the eye’s optics. Yet the neural visual system also adjusts to image blur (Fig. 9.11). Webster *et al.* (2002a) measured blur adaptation using methods similar to the ones described earlier to examine face adaptation. Images were blurred or sharpened by “distorting” the ratio of low to high frequency content in the images (or specifically, by varying the slope of the image amplitude spectra). Subjects adapted by viewing these biased images for a few minutes, and then adjusted the frequency spectrum of target images until they appeared properly focused. After viewing a blurry or sharpened image, a physically focused image appeared too sharp or too blurry, respectively (Fig. 9.12). Thus, the point of subjective focus shifted toward the

260 FITTING THE MIND TO THE WORLD




FIGURE 9.11 An illustration of visual adjustments to image blur (after Webster *et al.* 2002a). The two eyes at the center of each array are identical and physically focused, yet the right eye appears blurrier because it is embedded in a sharpened surround.




FIGURE 9.12 Adaptation to blur. Points plot the slope of the image amplitude spectrum that appeared best-focused after adapting to images with different spectral slopes. Steeper slopes (negative values) reduce the amplitude of higher frequencies more than lower frequencies and thus blur the images. Shallower slopes (positive values) boost the higher frequencies more than lower frequencies and thus sharpen the image. The four curves are for four different images. Prior adaptation to blurry images causes the original image to appear too sharp, and thus shifts the point of subjective focus toward the blurry adapting stimulus. Sharpened adapting images induce the opposite aftereffect.

adapting image. This also appears to happen in natural settings as clinicians seldom correct large refractive errors at once. Instead, patients with severe defocus or astigmatism prefer to have their optical aberrations corrected incrementally over time. Once again, these effects are consistent with a renormalization of perception, so that the currently viewed stimulus becomes the new prototype for proper image focus.

Like color, we can appreciate these effects in our own experience by considering the large changes in sensitivity across the visual field. Spatial resolution falls precipitously in the periphery, yet we do not experience the peripheral world as blurred (Galvin *et al.* 1997). Adaptation could again perceptually discount the optical and physiological variations because the responses at each retinal region will be normalized for the same external stimulus.

With regard to the themes we have been emphasizing, there are two important aspects in the compensations for optical differences between observers. The first is that we have better clues in the case of color or blur about the “priors” that the visual system is adjusting to (at least relative to the priors involved in processes like face recognition). As noted above, a number of studies have shown that the spatial statistics of natural images have a characteristic property (e.g. Burton & Moorhead 1987; Field 1987; Tolhurst *et al.* 1992; van der Schaaf & van Hateren 1996). Specifically, the amplitude spectrum of images falls inversely with spatial frequency (or as $1/f$), or in other words has a strong low-frequency bias. This scaling property of images is a common (though not universal) property of the physical world, and spatial sensitivity of the visual system is matched to it in many ways. A convincing perceptual demonstration of this can be seen when we look at filtered noise patterns. As Field and Brady (1997) noted, $1/f$ noise has salient structure at all spatial scales, while white noise, for which the amplitude is physically equal at all scales, instead appears dominated by the high-frequency components (Fig. 9.13).

The second important aspect is that for the eye’s optics we have a good understanding of how individuals differ, and where we can be confident that




FIGURE 9.13 The match between visual sensitivity and the spatial statistics of images. The white noise image at the right has a flat amplitude spectrum but perceptually appears to contain information only at fine scales. The middle image has a $1/f$ spectrum characteristic of natural images and instead appears to have equal energy at all scales.

262 FITTING THE MIND TO THE WORLD

these individual differences are very large. Specifically, we know that blur in the retinal image can vary widely because of differences in refractive errors, and that spectral sensitivity can vary widely because of the differences in the density of the lens pigment. Yet, individuals are often unaware of the changes in their optics, and when they become aware these are experienced as a failure of acuity or color discrimination rather than as a subjective experience that the world has changed. Arguably, this is because adaptation adjusts neural responses to maintain the balance of cortical responses to color and space, thus compensating for the imbalances introduced in the retinal image.

9.6 Time Scales and Types of Adaptation

Our review has focused on relatively dynamic perceptual adaptations, but it is important to emphasize that these are only one part of a wide array of different mechanisms that serve to adapt different aspects of visual coding. Table 9.2 lists some commonly recognized forms of visual plasticity ordered roughly in terms of the time scales over which they operate. (Note that this is only a subset from a still broader notion of plasticity that might include such phenomena as attention and priming.) While these involve very different mechanisms they may share a common functional role in shaping visual responses to fit properties of the world. Obviously, the properties that each adjusts to must vary depending on the time over which the information is integrated, and characterizing the relevant scene statistics over these integration times could provide important clues about the operating states that the visual system is trying to achieve. Yet the question of time scales is an aspect of adaptation that still remains very poorly defined. Some of the response changes associated with adaptation are extremely rapid, changing sensitivity with each brief fixation (Muller *et al.* 1999), and often only a few seconds of viewing are sufficient to induce powerful visual after-effects. However, there is also evidence for much slower perceptual adjustments.

Delahunt *et al.* (2004) examined the changes in color appearance following lens replacement in cataract patients. As pointed out above, the cataractous lens

Table 9.2 Mechanisms for perceptual plasticity and the time scales over which they may operate

Adjustment	Time scale
Evolution	Multiple life spans
Development	Months to years
Long term light and pattern adaptation	Minutes to months
Short term light and pattern adaptation	Milliseconds to minutes
Gain control	Milliseconds
Perceptual learning	?

AQ: Please provide the time scale for perceptual learning.




FIGURE 9.14 Changes in the white point following cataract surgery. Points plot the change in the achromatic settings (as distance in the standard CIE chromaticity diagram) for 4 observers as a function of post-operative days. Bold curve is the best fitting exponential function.

strongly filters short wavelength light, and thus immediately after the operation many patients report that the world appears bluish. Delahunt *et al.* tracked the changes in this perception by measuring the achromatic point at different times after surgery. Notably, white points in these subjects took weeks to return to values close to their pre-surgery settings (Fig. 9.14), even though changes in chromatic sensitivity were almost immediate. Long-term changes in color appearance have also been reported following exposure to biased light environments (Neitz *et al.* 2002). Neitz *et al.* exposed subjects to colored or filtered illuminants for several hours at a time, and found that this led to biases in perceived hue that could last for several days after observers returned to unbiased lighting. There are also reports of very long-term effects on spatial vision in patients with chronic cataracts. Fine *et al.* (2002) examined blur perception in an individual who had cataracts most of his life. After their removal, edges to him appeared to be too sharp, and this “aftereffect” showed little sign of diminishing even after months, suggesting that it might reflect a calibration fixed during development.

Webster *et al.* (2004) tested for analogous long-term adjustments in the perception of faces. They formed a series of images by morphing between a Japanese and a Caucasian face, and then asked subjects to select the boundary between the two ethnicities (i.e. the point at which the blended morph stopped looking Japanese and started looking Caucasian). These neutral points turned out to be very different for Asian and Caucasian observers, with each choosing a boundary closer to their own ethnic category (Fig. 9.15). Webster *et al.* used this difference to compare whether the settings might shift when individuals entered a new environment. For this, they compared the neutral points for Asian students who had either just arrived in the US or who had been resident for a year or more.

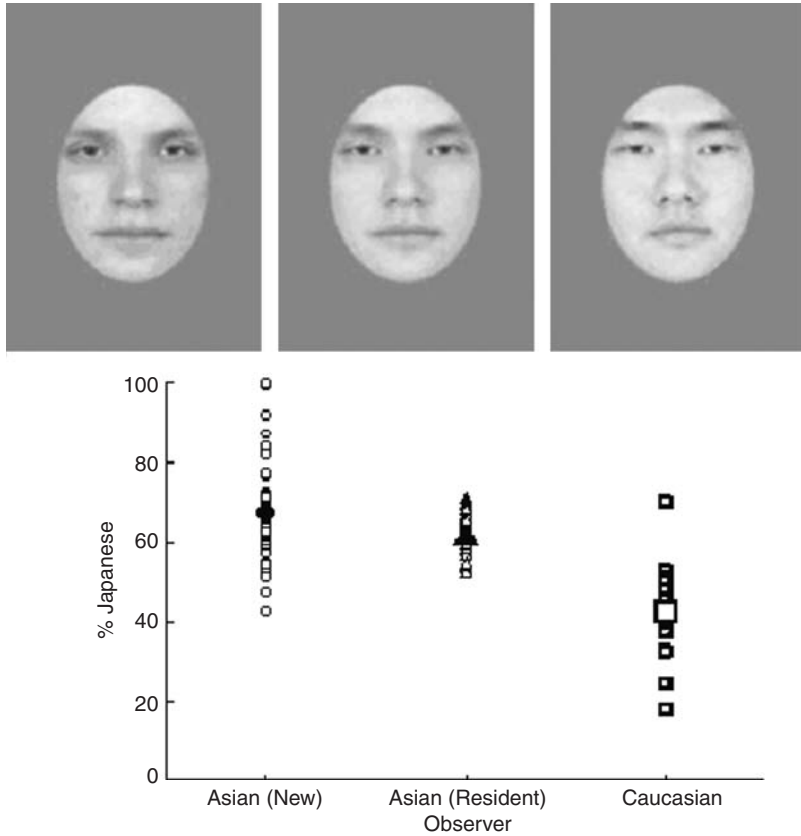


FIGURE 9.15 Changes in face perception. Each point shows the ethnicity boundary chosen by an individual subject in morphs between a Caucasian and Japanese face, as illustrated at the top by a Caucasian-Japanese face pair and an intermediate morphed image between them. Caucasian students in the US (points at right) chose a boundary closer to the Caucasian face, as if this face were more neutral to them. Newly arrived Asian students showed the opposite bias (points at left). Asian students who were resident in the US for a year or more show weak but significant shifts toward the Caucasian students' boundaries (middle points).

The latter group had a mean ethnicity boundary that was weakly but significantly shifted toward the mean settings for the Caucasian students. Moreover, these shifts in the direction of the Caucasian settings were positively correlated with their self-reports of the length of time in the US and with the proportion of time they interacted with Caucasian individuals while in the US. Each of these results is consistent with a gradual renormalization to their new face environment.

In order to probe these long-term adjustments, it is important to measure the neutral points in the absence of the relevant adapting stimulus. That is, to track

the patient's achromatic locus the testing is best done while the subject is dark adapted. What if the cataract patient instead first stares at a colored field, or if a Japanese observer stares at a Caucasian face? In these cases there is a very rapid adjustment to the current stimulus. For example, as shown before, viewing one face from a pair quickly and strongly biases the perceptual boundary in the morph between them, and a chromatically biased field quickly appears less saturated. Thus, when these stimuli are presented they rapidly capture the visual system and reset their sensitivity, but their effects also dissipate rapidly when the adapting stimulus is removed. There are other adaptation effects which have a much longer persistence. For example, the McCollough Effect can last for days, and this has suggested that the aftereffects do not passively decay but instead, are actively extinguished by the exposure to a different stimulus (Stromeyer 1978). However, the problem remains that for at least some adaptive adjustments there appear to be at least two distinct rates of sensitivity change, and possibly many. One represents an "extrinsic" state of adaptation that is controlled by the stimuli currently in front of us. However, underlying this is a more slowly drifting "intrinsic" state that integrates over a much longer time frame and thus over a much larger ensemble of stimuli. Very little is currently known about the dynamics and function of these long-term adjustments.

It is interesting to speculate that the extrinsic and intrinsic states might reflect adjustments to different aspects of visual coding. For example, most models of visual aftereffects depend on the notion of labeled lines – that the activity in a channel is associated with specific perceptual qualities. The aftereffects arise because adaptation selectively alters the distribution of responses across the channels (Webster 2003). However, it is possible that longer-term adjustments might involve changes in the *interpretation* of these responses, or in other words, in the labels the channels carry, and thus might reflect forms of "learning" that are very distinct from short-term adaptation. For example, one way to adapt to color would be to rescale the cone sensitivities so that the stimulus satisfies the "fixed" rule for white changes (e.g. so that the stimulus that equates responses across the three cone types changes). Another way is to change the rule, for instance by changing how the cone signals might be weighted within post-receptoral channels (Neitz *et al.* 2002).

9.7 Limits to Adaptation

Can we adapt to anything, so that aftereffects tell us only about the general malleability of perception, or does the presence of a specific aftereffect reveal something specific about the underlying skeletal structure of the visual system? Negative results have a tendency to go unreported, but as the number of new aftereffects continues to grow it may become increasingly interesting to identify situations in which the visual system fails to adapt. Such cases are important both for understanding which aspects of the environment are shaping our perception, and for understanding the extent to which adaptation can be used as a tool for characterizing the visual mechanisms.

One apparent limit is suggested by the finite selectivity of adaptation aftereffects. Exposure to a particular stimulus reduces sensitivity not only to that stimulus but also to similar ones, and thus this restricts the ways in which perception could be altered by adaptation. For example, in chromatic adaptation, the sensitivity changes are largely limited to independent response changes within a small number of chromatic mechanisms (“pi mechanisms”), which have spectral sensitivities that are similar (though not identical) to that of the photoreceptors (Stiles 1959). Similarly, after viewing a vertical grating, thresholds are elevated for a range of orientations around vertical, and these tuning curves have a bandwidth that is at least in rough correspondence with the average orientation selectivity of V1 cells (Blakemore & Campbell 1969). The limited selectivity of adaptation effects suggest that it should be possible to define “adaptation metamers,” or physically different stimuli that induce identical aftereffects, even though the two stimuli might be visually distinct. Such metamers should be specific for particular properties of the environment and to the particular visual mechanisms that adjust to those properties. As such, they provide a tool for isolating different stages of adaptation. For example, equivalent states of light adaptation can be induced by a steady field or a flickering field of the same mean luminance, since the property that light adaptation is, adjusting to is effectively only the time-averaged mean (Webster & Wilson 2000). However, these two stimuli look very different, and lead to very different adaptation effects in mechanisms that respond to luminance contrast, or to how light levels are varying around the mean. Thus the response changes of contrast adaptation are isolated by holding the state of light adaptation constant. Indeed, in studies of spatial pattern-adaptation, it is a routine procedure to move the stimulus over the retina in order to eliminate afterimages owing to local differences in light adaptation, and this is done in the hope of teasing out response changes that are specific to contrast adaptation. It is possible that similar techniques could be used to form pattern or contrast metamers to isolate still higher stages of adaptation. For example, it is possible in principle to specify different sets of colors that are perceptually distinct but result in equivalent changes of color contrast adaptation (e.g. identical aftereffects to those shown in Fig. 9.6). However, because these are perceptually different, there may be other ways in which the visual system adapts differently to the two distributions (e.g. changing how noticeable or salient different colors are after adapting to either distribution), and thus such pairs might isolate adaptive adjustments that are more closely tied to the phenomenal appearance of the stimuli.

A further possible constraint concerns the nature of contingent adaptation. As we noted before, the McCollough Effect (ME) is usually interpreted as a color afterimage that is contingent on orientation. A number of studies have explored whether analogous aftereffects can be induced by pairing color with more arbitrary spatial patterns. Strong ME’s have been reported for a variety of geometrical patterns (e.g. concentric rings versus radial spokes) (Humphrey *et al.* 1985). Conversely, spatial aftereffects can be made contingent on color. An intriguing case was reported by Durgin (1996), in which texture density

aftereffects were selective for the color of a surrounding frame. Such results suggest that the visual system can adapt to a potentially wide range of associations, and there are arguments that contingent aftereffects may actually represent a form of learning rather than sensory adaptation (Siegel & Allan 1992). However, it remains controversial whether a ME can be formed with any pattern. For example, reports of ME's for different printed words have been difficult to replicate (Humphrey *et al.* 1994). Yamashita *et al.* (2000) tested for ME's when red and green were paired with a contracted or expanded face. No aftereffects were visible, even though the spatial differences between the two adapting patterns were very obvious. Thus the ME's that can be formed may say something non-arbitrary about visual coding.

A final important but largely unexplored limit to adaptation may be individual differences in the degree or form of the adaptation. We have assumed for the present arguments that adaptation is uniform across different individuals, but to the extent that two observers differ in their adaptability, they should be held in phenomenally different worlds even within the same environment. A particularly intriguing case to consider is adaptation in the aging visual system. As noted above, adaptation may play a crucial role in maintaining perceptual constancy despite the large optical and neural changes that accompany normal aging. Surprisingly, very little is known about how the processes of adaptation – and in particular how cortical processes of pattern-selective adaptation – might themselves age. Deficits in adaptation could have significant and unfavorable consequences for perception, for as we have seen, discrimination is best around the stimulus levels we are adapted to. Failures to adjust this level to the appropriate environment might hinder perceptual distinctions (e.g. leaving the individual with a perpetual “other-race effect”), and could be visually discomforting by inflating stimulus salience (e.g. if the world always appeared noticeably blurred). Studies of adaptation in the aging visual system might also shed light on the underlying mechanisms of the response changes. For example, to the extent that adaptation reflects passive fatigue, then factors that compromise visual function might result in more pronounced adaptation. Alternatively, if the sensitivity changes depend on an active mechanism for calibrating neural sensitivity, then factors that weaken efficiency should reduce adaptation.

9.8 Adaptation and the Contents of Visual Awareness

The preceding sections explore the possibility that the states of adaptation vary in more or less predictable ways, and that these variations constrain whether two different observers have similar or different subjective experiences. In this section, we consider how adaptation might influence the actual content of their experience. Visual aftereffects are among the most striking phenomena in perception, and can quickly alter the appearance of the world in dramatic and startling ways – the sensation of a contracting world after watching an expanding spiral arouses the curiosity of even the most apathetic observer. Yet it has thus far

proven surprisingly difficult to demonstrate that these perceptual shifts are accompanied by equally dramatic shifts in visual performance. Most performance measures have focused on the ability to discriminate changes in the stimulus following adaptation (e.g. detecting a difference in contrast or orientation). Better discrimination is predicted if adaptation shifts the response curves of visual mechanisms so that they are steepest and thus most sensitive around the adapting level. This clearly occurs in the case of light adaptation, where it is easy to show that we can discriminate the contrasts in a scene better after we adjust to the average brightness (Whittle 1992); and in individual neurons, which have limited dynamic range for intensity and contrast (Ohzawa *et al.* 1985). But whether similar improvements in pattern discrimination occur following adaptation to patterns remains controversial, and even where these have been reported, they seem weak compared to the sheer magnitude of the perceptual shifts (e.g. Regan & Beverley 1985; Greenlee & Heitger 1988; see also the chapter by Rhodes *et al.*, this volume). This suggests that some of the principal consequences and functions of adaptation may lie in how they influence phenomenology. The startle and awe that visual aftereffects engender may turn out to be the best measure of their utility.

9.8.1 Perceptual Norms

Adaptation produces a number of interrelated effects on phenomenology. One set of effects concerns how we experience the adapting stimulus itself, while a second set concerns our perception of new stimuli, to which we are not adapted. As we have stressed throughout, one of the chief consequences of adaptation is to normalize visual coding relative to the stimuli that we have been exposed to. Through this adjustment, the adapting stimuli define the neutral points of visual coding. But there is also a strong sense in which they become the neutral points of visual experience. The white point has a special status in color perception because it provides the chromatically neutral reference point or norm against which other stimuli can be contrasted. It represents the average or expected color percept, yet again what appears gray is to a large extent the average of the set of colors that we have encountered.

Calibrating the perception of different observers relative to a common external reference may prove to be an essential precursor to the ability to communicate about our perceptions with others, and could underlie some important aspects of norms and aesthetics within a population that might instead be attributed to a common culture. That is, these collective norms may reflect shared perceptions rather than shared criteria. Judgments of facial attractiveness is an example of this. A number of studies have found that faces that are physically average tend to be rated as more attractive than more distinctive faces (Langlois & Roggman 1990). However, we are all exposed to a different distribution of faces, and thus the average that is important to any individual may be the average of the particular distribution they have seen. Rhodes *et al.* (2003) recently showed that exposure to a set of distorted faces biased the rated attractiveness of a different

set of test faces. These shifts are consistent with renormalizing face coding with respect to the adapting faces, so that these faces appeared less distorted and thus, more attractive. This could also conceivably explain the sometimes-striking tendency of spouses and partners to look alike. They may be visually attracted to individuals that match the facial configurations to which they are adapted, for example because of exposure to a particular community or their relatives or themselves. Perceptual adjustments could also provide a simple sensory explanation of how judgments of attractiveness might change when an individual is exposed to a new visual environment, as the following description of the anthropologist Malinowski's experiences suggests:

AQ: This is not mentioned in the reference list.

“Malinowski (1929) makes the intriguing observation that after he had lived in the Trobriand Islands for some time his judgments of Trobriand beauty began to agree with the Trobrianders' judgments.” *Symons*, pg 196 (Symons 1979).

Within the present framework, it was not so much Malinowski's perception of beauty that changed. Rather, what changed was the physical stimulus that induced that perception, through the changed state of adaptation to his new environment.

It is important to note that perceptual norms may be important not only for judging the average stimulus, but also for categorizing stimuli. For example, in the perception of a facial quality like gender, images corresponding to the mean may be rare. Yet adaptation to the population mean may nevertheless be important for defining the category boundary. It is possible that in contexts where the distribution of males and females is highly skewed (e.g. in some institutionalized settings), judgments of gender would become perceptually skewed because of adaptation-induced shifts in the category boundaries.

9.8.2 Perceptual Gamuts

A second consequence of renormalization may be to expand perceptual space in order to match it to the range of levels in the stimulus. This is similar to the notion of sphering in visual coding, so that equal perceptual weight is given to the variance along different stimulus dimensions. Distorting the range of perceptual responses may not affect the ability to discriminate between different stimulus levels, for simply boosting the gain of a mechanism will not improve sensitivity if it increases both signal and noise. Thus, the impact of these adjustments on perception may not be readily captured in measures of visual sensitivity. An interesting example is provided by the perceptual experience of anomalous trichromats. The difference in spectral sensitivity between their L and M cone pigments is much smaller than in color normals, and thus their color discrimination is much poorer. Yet they may not “experience” the world as desaturated because this signal is amplified in postreceptoral channels (MacLeod 2003). In the same way, we noted previously that in normal observers the signals from S cones are amplified, and this perceptually spheres color space even though it does not undo sensitivity limits imposed by the paucity of S cones.




FIGURE 9.16 Perceptual distortions of stimulus space. Top row: three examples of white noise. Though these images are statistically very different they appear very similar and are difficult to discriminate. Middle row: three examples of foliage. Natural images share a great deal of common statistical structure, but their differences are much more readily perceived. Bottom row: three examples of faces (devotees at the Nashik Kumbh Mela). Faces are highly similar physically, but differences between them are more salient than most other image classes.

Figure 9.16 is a further possible illustration of these perceptual distortions. The top row shows three examples of white noise, in which the value of each pixel is chosen at random. These images are physically quite different. That is, their RMS difference (i.e., their Euclidean distance) is quite large. However, perceptually they appear quite close. In contrast, the three middle images of outdoor scenes are conspicuously different, yet their statistics are highly constrained. In fact, the set of all natural images fills only a tiny fraction of the state space of all possible images, yet the visual system distorts this space by expanding the response to natural images (and by perceptually compressing the responses to images that fall outside the natural gamut) (Field 1994). In the same way, we can find distortions across different classes of natural images. For example, the bottom row shows three examples of faces. The set of possible

faces again occupies only a miniscule volume within the space of natural images, yet the differences between faces are often much more conspicuous and much easier to remember than for other classes of objects. These perceptual differences may reflect both evolutionary and short-term adaptations that amplify visual responses to the characteristic properties of the visual environment.

The principle of perceptual expansion can be applied at a still finer level by comparing visual responses within the category of face images. An example is the “other race” effect, or the tendency to easily distinguish among the kinds of faces we are used to seeing while finding it difficult to distinguish among individuals from other groups (Furl *et al.* 2002). If adaptation matches face coding to the specific gamut of faces we encounter, then this will serve to highlight the differences between faces around the average. This predicts that for physically equivalent stimulus differences, judgments of phenomenal similarity would increase for images that are far removed from the average face, a behavior consistent with the other-race effect.

What are the advantages of biasing phenomenal differences if they do not lead to comparable improvements in visual discrimination? The answer may be that they do help discriminations, but not of the type that are usually measured. The classic work of Miller (Miller 1956) illustrated that perceptual judgments are limited to a finite number of categories. Adjusting the response gamut might aid perception by making the full range of categories available.

9.8.3 *Perceptual Salience*

A second way in which adaptation could aid perception is by highlighting how a stimulus differs from the stimuli we are adapted to. Barlow (1990) proposed that adaptation reflects a form of learning about the associations underlying the structure of the visual world. This is important for coding efficiency because adjusting to these associations allows us to encode the world within mechanisms whose signals vary independently. But another effect of this is to draw attention to new associations, or “suspicious coincidences” in the world. That is, according to Barlow, adaptation is a process that brings new properties of the environment to our notice. Note that these novel properties are the very stimuli we experience in visual aftereffects. In a red world a flat spectrum stands out as green, and in a world of expanded faces a neutral face will appear conspicuously contracted. Note also that when we experience an adaptation effect, it is the aftereffect that strikes us much more than any perceptual changes in the adapting stimulus itself. For example, when we stare at a waterfall we are largely unaware of the changes in motion sensitivity, but these changes are overwhelming when we switch our gaze to the novel properties of the surrounding static scene. Yet we are always adapted, and thus we are always experiencing an adaptation effect. Thus, much of the content that reaches our conscious awareness may be a visual aftereffect.

Another way to conceptualize these perceptual consequences is from the perspective of predictive coding (Srinivasan *et al.* 1982; Mumford 1994).

The current state of adaptation is in a sense a prediction about the current state of the world. This prediction may be represented only implicitly in our conscious awareness, for it forms a template that filters out the expected stimulus structure and thus passes only the errors. These errors represent the novel properties of the image, and these become accessible only by first discounting the properties that are old. One implication of this is that as we adapt to a new environment we should become better at detecting novel features within it. Consistent with this, Webster *et al.* (1998) found in a visual search task that observers were faster at finding a colored target if they were first adapted to the colored backgrounds on which they had to search. Similarly, this would predict that in “change blindness” tasks, which measure the failures to detect alterations in scenes, (Rensink *et al.* 1997) that the changes observers do or do not notice should depend in part on whether or not they are adapted to properties of the scenes.

The role of adaptation in modulating visual salience suggests that there is a close relation between adaptation and attention, and there are many examples of interactions between them (Boynton 2004). Other chapters in this book discuss the role of attention in modulating the magnitude of perceptual adaptation. Here we note that, conversely, the state of adaptation may strongly modulate what aspects of the world we attend to.

With these thoughts in mind, we can now return to Freud’s portraits. The very act of exploring the face as he created it may have strongly adapted him to it, so that he perceived it – visually – in a way that perhaps no one else can. While at first glance others are struck by salient distortions and stylistic similarities, in his eye the same stimulus properties might be dulled by a process that calibrates normal only according to the history of stimulation. This is in fact a common anecdotal impression. Faces that appear striking to us when we first see them typically lose their distinctiveness over time, and this can be parsimoniously accounted for by actual shifts in perception. Our analysis of Freud’s paintings is not intended as an aesthetic judgment about the work, for we instead suggest that differences in aesthetics may have roots in perceptual differences. Nor do we mean to preclude the possibility that the artist chose to emphasize or perhaps exaggerate certain traits that are perceptually obvious to him. We used the work merely as an example to say that, if we could look at the painting through his eyes, it would not look the same.

9.9 Measuring Conscious Experience

The central problem in the phenomenology of perception is that it is a private experience, and thus we have access only to our own. In this chapter, we have argued that this private experience is shaped in important ways by processes of adaptation, and have assumed that these physiological processes are similar in important ways within different individuals (in that the common effect of adaptation is to normalize neural activity). We have also argued that we are always

adapted to specific properties of our environment, and that it is these properties that ultimately control the states of adaptation. Thus, some aspects of our inner private experience depend on outer public variables. To measure what world an observer would experience as physically focused (i.e. not blurred), we may not need to measure the observer. We can instead measure the spatial statistics of their environment. To ask whether you and I have different experiences, we can measure whether our environments differ. There is a rapidly growing interest in characterizing natural scene statistics because they hold the promise of revealing much about visual coding. Ultimately, these measurements may also reveal much about visual experience.

9.10 Epilogue

At the time of this writing, news was focused on another royal portrait, this time of Prince Phillip. Commissioned by the Royal Society for the Encouragement of the Arts, Manufactures and Commerce, of which Phillip is President, the portrait was painted by the artist Stuart Pearson Wright, and showed a bare-chested Phillip with cress growing out of his finger (a strand for each child) and a bluebottle fly resting on his shoulder (to signify mortality). Phillip refused further sittings after seeing the work in progress. The painting was rejected by the RSA, and Wright instead provided a second image showing the Prince only from the neck up, but the controversy recently resurfaced when the original was put on the market. We note it here because it again bears on the theme of whether different individuals see the world in similar ways. Asked by Wright if he had caught his likeness, the Prince replied "I bloody well hope not!"

Acknowledgments

Supported by the National Eye Institute (EY10834), the National Institute on Aging (AG04058), and a Jules and Doris Stein RPB Professorship to JSW. Parts of this chapter are based on Webster (2002).

References

- Andrews, D.P. (1964). Error-correcting perceptual mechanisms. *Quarterly Journal of Experimental Psychology*, 16, 104–15.
- Anstis, S.M. (1996). Was El Greco astigmatic? *Investigative Ophthalmology and Visual Science*, 37, S697.
- Atick, J.J. (1992). Could information theory provide an ecological theory of sensory processing? *Network*, 3, 213–51.
- Barlow, H.B. (1990). A theory about the functional role and synaptic mechanism of visual aftereffects. In C. Blakemore (Ed.), *Vision: Coding and Efficiency* (pp. 363–75). Cambridge: Cambridge University Press.

- Barlow, H.B., & Földiák, P. (1989). Adaptation and decorrelation in the cortex. In R. Durbin, C. Miall, & G.J. Mitchison (Eds.), *The Computing Neuron* (pp. 54–72). Wokingham: Addison-Wesley.
- Barlow, H.B., & Levick, W.R. (1976). Threshold setting by the surround of cat retinal ganglion cells. *Journal of Physiology, London B*, 212, 1–34.
- Blakemore, C., & Campbell, F.W. (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. *Journal of Physiology*, 203, 237–60.
- Boynton, G.M. (2004). Adaptation and attentional selection. *Nature Neuroscience*, 7, 8–10.
- Brainard, D.H., & Wandell, B.A. (1992). Asymmetric color matching: how color appearance depends on the illuminant. *Journal of the Optical Society of America A*, 9, 1433–48.
- Brown, R.O., & MacLeod, D.I.A. (1997). Color appearance depends on the variance of surround colors. *Current Biology*, 7, 844–9.
- Buchsbaum, G., & Gottschalk, A. (1983). Trichromacy, opponent colours and optimum colour information transmission in the retina. *Proceedings Royal Society London B*, 220, 89–113.
- Burton, G.J., & Moorhead, I.R. (1987). Color and spatial structure in natural scenes. *Applied Optics*, 26, 157–70.
- Carandini, M., Barlow, H.B., O’Keefe, L.P., Poirson, A.B., & Movshon, J.A. (1997). Adaptation to contingencies in macaque primary visual cortex. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 52, 1149–54.
- Chaparro, A., Stromeyer, C.F.I., Huang, E.P., Kronauer, R.E., & Eskew, R.T.J. (1993). Colour is what the eye sees best. *Nature*, 361, 348–50.
- De Valois, R.L. (2003). Neural coding of color. In L.M. Chalnpa, & J.S. Werner (Eds.), *The Visual Neurosciences Volume 2*. (pp. 1003–16). Cambridge: MIT Press.
- De Valois, R.L., Cottaris, N.P., Elfar, S.D., Mahon, L.D., & Wilson, J.A. (2000). Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proceedings of the National Academy of Science*, 97, 4997–5002.
- Delahunt, P., Webster, M.A., Ma, L., & Werner, J.S. (2004). Long-term renormalization of chromatic mechanisms following cataract surgery. *Visual Neuroscience*, in press.
- Dragoi, V., Sharma, J., Miller, E.K., & Sur, M. (2002). Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nature Neuroscience*, 5, 883–91.
- Durgin, F.H. (1996). Visual aftereffect of texture density contingent on color of frame. *Perception and Psychophysics*, 58, 207–23.
- Field, D.J., & Brady, N. (1997). Visual sensitivity, blur, and the sources of variability in the amplitude spectra of natural images. *Vision Research*, 37, 3367–83.
- Field, D.J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4, 2379–94.
- Field, D.J. (1994). What is the goal of sensory coding? *Neural Computation*, 6, 559–601.
- Fine, I., Smallman, H.S., Doyle, P., & MacLeod, D.I.A. (2002). Visual function before and after removal of bilateral congenital cataracts in adulthood. *Vision Research*, 42, 191–210.
- Fiorentini, A., Baumgartner, G., Magnussen, S., Schiller, P.H. & Thomas, J.P. (1990). The perception of brightness and darkness: Relations to neuronal receptive fields. In L. Spillmann & J.S. Werner (Eds.), *Visual Perception: The Neurophysiological Foundations* (pp. 129–61). San Diego: Academic Press.

AQ: This reference is not cited in the text.

- Furl, N., Phillips, P.J., & O'Toole, A.J. (2002). Face recognition algorithms and the other-race effect: computational mechanisms for a developmental contact hypothesis. *Cognitive Science*, 26, 797–815.
- Galvin, S.J., O'Shea, R.P., Squire, A.M., & Govan, D.G. (1997). Sharpness overconstancy in peripheral vision. *Vision Research*, 37, 2035–9.
- Golz, J., & MacLeod, D.I.A. (2002). Influence of scene statistics on colour constancy. *Nature*, 415, 637–40.
- Greenlee, M.W., & Heitger, F. (1988). The functional role of contrast adaptation. *Vision Research*, 28, 791–7.
- Hardin, C.L. (1997). Reverting the spectrum. In A. Byrne & D. Hibert (Eds.), *Readings on Color, Volume 1: The Philosophy of Color* (pp. 289–301). Cambridge: MIT Press.
- Hardy, J.L., Frederick, C.M., Kay, P., & Werner, J.S. (2005). Color naming, lens aging, and grue: What the optics of the aging eye can teach us about color language. *Psychological Science*, in press.
- Hecht, S. (1923). Sensory adaptation and the stationary state. *Journal of General Physiology*, 5, 555–79.
- Helson, H. (1964). *Adaptation-Level Theory*. New York: Harper and Row.
- Humphrey, G.K., Dodwell, P.C., & Emerson, V.F. (1985). The roles of pattern orthogonality and color contrast in the generation of pattern-contingent color aftereffects. *Perception and Psychophysics*, 38, 343–53.
- Humphrey, G.K., Skowbo, D., Symons, L.A., Herbert, A.M., & Grant, C.L. (1994). Text-contingent color aftereffects: a reexamination. *Perception and Psychophysics*, 56, 405–13.
- Hurvich, L.M., & Jameson, D. (1974). Opponent processes as a model of neural organization. *American Psychologist*, 29, 88–102.
- Langlois, J.H., & Roggman, L.A. (1990). Attractive faces are only average. *Psychological Science*, 1, 115–21.
- Laughlin, S.B. (1987). Form and function in retinal processing. *Trends in Neuroscience*, 10, 478–83.
- Lennie, P. (1999). Color coding in the cortex. In K. Gegenfurtner & L.T. Sharpe (Eds.), *Color Vision: From Genes to Perception* (pp. 235–47). Cambridge: Cambridge University Press.
- Lennie, P., & D'Zmura, M. (1988). Mechanisms of color vision. *CRC Critical Reviews of Neurobiology*, 3, 333–400.
- Lindsey, D.T., & Brown, A.M. (2002). Color naming and the phototoxic effects of sunlight on the eye. *Psychological Science*, 13, 506–12.
- Locke, J. (1689, 1975). *An Essay Concerning Human Understanding*. Oxford: Clarendon Press.
- MacLeod, D.I.A. (2003). Colour discrimination, colour constancy, and natural scene statistics. In J.D. Mollon, J. Polorny, & K. Kroblanch (Eds.), *Normal and Defective Colour Vision*. Oxford: Oxford University Press.
- MacLeod, D.I.A., & von der Twer, T. (2003). The pleistochrome: optimal opponent codes for natural colors. In R. Mausfeld & D. Heyer (Eds.), *Colour Perception: Mind and the Physical World* (pp. 155–84). Oxford: Oxford University Press.
- McCollough, C. (1965). Color adaptation of edge-detectors in the human visual system. *Science*, 149, 1115–6.
- Miller, G.A. (1956). The magic number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.

AQ: Please clarify the year of publication.

276 FITTING THE MIND TO THE WORLD

- Mizokami, Y., Webster, M.A., & Webster, S.M. (2003). Seasonal variations in the color statistics of natural images. *Journal of Vision*, 3, 444.
- Movshon, J.A., & Lennie, P. (1979). Pattern-selective adaptation in visual cortical neurones. *Nature*, 278, 850–2.
- Muller, J.R., Metha, A.B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science*, 285, 1405–8.
- Mumford, D. (1994). Neuronal architectures for pattern theoretic problems. In C. Koch & J.L. Davis (Eds.), *Large Scale Neuronal Theories of the Brain* (pp. 125–52). Cambridge: MIT Press.
- Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M., & Williams, D.R. (2002). Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron*, 35, 783–92.
- Ohzawa, I., Sclar, G., & Freeman, R.D. (1985). Contrast gain control in the cat's visual system. *Journal of Neurophysiology*, 54, 651–67.
- Olshausen, B.A., & Field, D.J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–9.
- Palmer, S.E. (1999). Color, consciousness, and the isomorphism constraint. *Behavioral and Brain Sciences*, 22, 923–89.
- Párraga, C.A., Brelstaff, G., Troscianko, T., & Moorehead, I.R. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America A*, 15, 563–9.
- Regan, D., & Beverley, K. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America A*, 2, 147–55.
- Rensink, R.A., O'Regan, J.K., & Clark, J.J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science*, 8, 368–73.
- Rhodes, G., Jeffery, L., Watson, T.L., Clifford, C.W.G., & Nakayama, K. (2003). Fitting the mind to the world: face adaptation and attractiveness aftereffects. *Psychological Science*, 14, 558–66.
- Ruderman, D.L., Cronin, T.W., & Chiao, C-C. (1998). Statistics of cone responses to natural images: implications for visual coding. *Journal of the Optical Society of America A*, 15, 2036–45.
- Siegel, S., & Allan, L.G. (1992). Pairings in learning and perception: Pavlovian conditioning and contingent aftereffects. In D. Medin (Ed.), *The Psychology of Learning and Motivation* (Vol. 28, pp. 127–60). New York: Academic Press.
- Simoncelli, E.P., & Olshausen, B.A. (2001). Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24, 1193–216.
- Srinivasan, M.V., Laughlin, S.B., & Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London B*, 216, 427–59.
- Stiles, W.S. (1959). Color vision: the approach through increment-threshold sensitivity. *Proceedings of the National Academy of Sciences*, 45, 100–14.
- Stromeyer, C.F.I. (1978). Form-color aftereffects in human vision. In R. Held, H.W. Leibowitz, & H.L. Terrber (Eds.), *Handbook of Sensory Physiology*, (Vol. VIII, pp. 97–142). New York: Springer-Verlag.
- Symons, D. (1979). *The Evolution of Human Sexuality*. New York: Oxford University Press.
- Tolhurst, D.J., Tadmor, Y., & Chao, T. (1992). Amplitude spectra of natural images. *Ophthalmology Physiological Optics*, 12, 229–32.
- van der Schaaf, A., & van Hateren, J.H. (1996). Modelling the power spectra of natural images: statistics and information. *Vision Research*, 36, 2759–70.
- von Kries, J. (1970). Chromatic adaptation, in Festschrift der Albrecht-Ludwigs-Universität (Fribourg, 1902). Trans. D. L. MacAdam (Ed.), *Sources of Color Science* (pp. 109–19). Cambridge: MIT Press.

- Wainwright, M.J. (1999). Visual adaptation as optimal information transmission. *Vision Research*, *39*, 3960–74.
- Walraven, J., Enroth-Cugell, C., Hood, D.C., MacLeod, D.I.A., & Schnapf, J.L. (1990). The control of visual sensitivity: receptor and postreceptor processes. In L. Spillman & J.S. Werner (Eds.), *Visual Perception The Neurophysiological Foundations* (pp. 53–101). San Diego: Academic Press.
- Webster, M.A. (2002). Adaptation, high-level vision, and the phenomenology of perception. In B.E. Rogowitz, & T. Pappas (Eds.), *Human Vision and Electronic Imaging VII SPIE*, 4662. (pp. 1–11).
- Webster, M.A. (2003). Pattern-selective adaptation in color and form perception. In L.M. Chalupa, & J.S. Werner (Eds.), *The Visual Neurosciences*, (Vol. 2, pp. 936–47). Cambridge: MIT Press.
- Webster, M.A., Georgeson, M.A., & Webster, S.M. (2002a). Neural adjustments to image blur. *Nature Neuroscience*, *5*, 839–40.
- Webster, M.A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, *428*, 558–61.
- Webster, M.A., Malkoc, G., Bilson, A.C., & Webster, S.M. (2002b). Color contrast and contextual influences on color appearance. *Journal of Vision*, *2*, 505–19.
- Webster, M.A., & Mollon, J.D. (1994). The influence of contrast adaptation on color appearance. *Vision Research*, *34*, 1993–2020.
- Webster, M.A., & Mollon, J.D. (1997). Adaptation and the color statistics of natural images. *Vision Research*, *37*, 3283–98.
- Webster, M.A., Raker, V.E., & Malkoc, G. (1998). Visual search and natural color distributions. In B.E. Rogowitz, & T. Pappas (Eds.), *Human Vision and Electronic Imaging III SPIE*, 3299 (pp. 498–509).
- Webster, M.A., Webster, S.M., Bharadwaj, S., Verma, R., Jaikumar, J., Madan, G., et al. (2002c). Variations in normal color vision III. Unique hues in Indian and United States observers. *Journal of the Optical Society of America A*, *19*, 1951–62.
- Webster, M.A., & Wilson, J.A. (2000). Interactions between chromatic adaptation and contrast adaptation in color appearance. *Vision Research*, *40*, 3801–16.
- Werner, J.S. (1998). Aging through the eyes of Monet. In W.G.K. Backhums, R. Kliegel, & J.S. Werner (Eds.), *Color Vision Perspectives from Different Disciplines* (pp. 3–41). Berlin: de Gruyter.
- Werner, J.S., & Scheffrin, B.E. (1993). Loci of achromatic points throughout the life span. *Journal of the Optical Society of America A*, *10*, 1509–16.
- Whittle, P. (1992). Brightness, discriminability, and the ‘Crispening Effect’. *Vision Research*, *32*, 1493–1507.
- Wright, W.D. (1928–29). A re-determination of the trichromatic coefficients of the spectral colors. *Trans. Opt. Soc. London*, *30*, 141–64.
- Wright, W.D. (1988). Talking about color. *Color Research and Application*, *13*, 138–9.
- Yamashita, J.A., Hardy, J.L., De Valois, K.K., & Webster, M.A. (2000). Color and contrast-selective after-effects in the perception of faces. *Invest. Ophthalmol. Suppl.* *41*, S224.
- Young, R.W. (1991). *Age-related Cataract*. New York: Oxford University Press.

AQ: Please expand this.

AQ: Please provide complete journal names of these two to incorporate in the reference list.

AQ: This reference is not cited in the text.

