

Adaptation, High-Level Vision, and the Phenomenology of Perception

Michael A. Webster*

Department of Psychology, University of Nevada, Reno

ABSTRACT

To what extent do we have shared or unique visual experiences? This paper examines how the answer to this question is constrained by known processes of visual adaptation. Adaptation constantly recalibrates visual sensitivity so that our vision is matched to the stimuli that we are currently exposed to. These processes normalize perception not only to low-level features in the image, but to high-level, biologically relevant properties of the visual world. They can therefore strongly impact many natural perceptual judgments. 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 will differ. These differences are illustrated by considering how adaptation can influence human face perception. 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 is illustrated by reviewing the effects of adaptation on the perception of image blur. In either case, it is the similarities or differences in the stimuli – and not the intrinsic similarities or differences in the observers – which 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.

1. INTRODUCTION

Recently a new portrait of Queen Elizabeth II was unveiled by the artist Lucian Freud. Freud (grandson of Sigmund) has been described as the greatest living artist in England, and he clearly labored long over a project that included 70 separate sittings by the Queen. Yet the work was not well received. Criticisms in the press ranged from mild disappointment to open hostility. Many of these pointed to distortions in the representation (to the stubble of a beard, to the neck of a rugby player, to the contorted expression of the Queen's dog following a stroke). Evidently, the public saw the painting in a way that the artist did not. In this paper I want to argue that they *literally saw* the painting differently. By this I am not suggesting that anyone perceived the portrait as a veridical depiction of the Queen, an error of logic known by the El Greco fallacy¹. Instead, I will explore the possibility that the painting looked different to Freud, simply because he had spent so much time looking at it.

A central philosophical question in perception is how the world might appear to us if we could see it through the eyes of another. The private nature of our internal experience may forever preclude a complete answer. Yet analyses of the relationships between perceptions (e.g. of the similarities between different colors) have been used to argue that visual experience has a similar structure across individuals^{2,3}. The present paper examines the implications of visual adaptation for conscious experience – and in particular for the question of whether two individuals have a common or distinct visual experience. Adaptation is a ubiquitous and intrinsic property of vision, adjusting sensitivity to many aspects of the stimuli we are currently exposed to⁴. These adjustments determine the operating range and response properties of the visual system, and thus profoundly affect the way things look. They could therefore be a fundamental factor determining whether things look the same or different to others.

2. LIGHT ADAPTATION

The perceptual consequences of adaptation can be illustrated by a simple example of light adaptation. Imagine two “identical” observers, one in a room illuminated by white light and the other sitting a room lit by red light (Figure 1). Because they are exposed to a different average color, light adaptation within the retina will adjust their vision in different ways. A common characteristic of these adjustments is that they serve to normalize responses across the set of mechanisms. In chromatic or von Kries adaptation, the normalization operates through a gain control in cone-specific

* mwebster@unr.nevada.edu; phone 775-784-6828 ext 2029; fax 775-784-1126; Department of Psychology/296, University of Nevada, Reno, Reno NV USA 89557

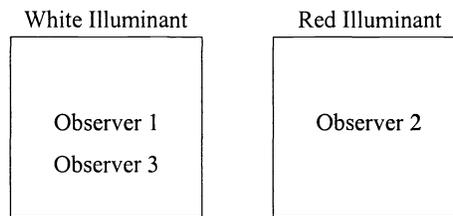


Figure 1. Color perception under different illuminants. Observers 1 and 2 have identical spectral sensitivities but will experience the same color stimuli in different ways, because they are adapted to different environments. Observers 1 and 3 have different spectral sensitivities but will experience color in common ways, because they are adapted to the same environment.

pathways, which keeps sensitivity inversely proportional to the strength of the input⁵. For example, in the red room the longwave-sensitive cones will absorb more quanta, and thus light adapt to become less sensitive, while the shortwave-sensitive cones will catch fewer quanta and thus dark adapt to become more sensitive. As a result, the average output of each mechanism (e.g. each cone receptor) remains constant across different viewing conditions, while the average responses of different mechanisms become equated for any given viewing condition. Taken to completion, the perceptual effects of these sensitivity changes would be to recalibrate vision so that in either room the average lighting appears gray.

But suppose they could suddenly see the rooms through each other's differently-adapted eyes. Clearly, their impressions would not agree. The observer from the "gray" room would see the second room as red, and thus would not agree with the other observer's impression that it was gray. Conversely, the observer adapted to the "red" room would find that the second room appeared green, and again would not agree with the other observer's experience of gray. The point then, is that two observers who are under different states of adaptation will experience the same stimulus in different ways.

What if we now place a third observer in the gray room, one who has a different visual system? For example, the observer might have a more pigmented lens so that she is less sensitive to shorter wavelengths. The light reaching her cones will differ, but over time the retina of this observer will once again adapt to the average spectrum in the room, equating the cone responses so that the illuminant appears gray. If our original observer could now look through her eyes they would agree in their experience of gray. Here then, the point is that two observers who are adapted to the same stimulus should – in important ways – experience the world in a common way. Note this is the case even though their visual systems are intrinsically different. In this sense, whether two individuals have shared or distinct perceptions depends more on properties of the environment than on properties of the individuals.

In fact, we do not need to imagine these effects in the minds of others, because we can experience such comparisons within ourselves. Many of the sensitivity changes in adaptation are specific to the region of the retina that was exposed to the stimulus. This is especially true of the peripheral adjustments of light adaptation. When we stare at a simple pattern like a red circle, the cones that are exposed to the image light adapt to the red, and thus migrate toward a different state of adaptation than the surrounding retina. With careful fixation the circle grows less saturated over time, and can even disappear entirely. Yet if we shift our view to the side, the paper that to part of our retina appeared white is now perceived as a vivid green afterimage. Similarly, if we examine a blank sheet of paper it appears uniform in color, even though the retina itself is highly inhomogeneous in the number and relative proportion of different receptors and postreceptoral neurons, and in the density of screening pigments. Despite these physiological differences the color of a stimulus remains remarkably stable regardless of its position in the visual field⁶, and it seems plausible to suppose that this results in part because physiologically different parts of the retina are in general adapted for the same average stimulus in the environment. Here again, it is the common feature of the environment – and not different parts of the observer - that induces a common and consistent visual experience. (However, color in the world is not completely isotropic. For example, the lower retina may more often be exposed to sky, and it would be of interest to examine whether long-term adaptation to such spatial nonuniformities of color in the environment induce corresponding nonuniformities in color perception across the visual field.)

It should be noted that O'Regan and Noe⁷ recently proposed a different account of the stability of color perception across the visual field. They suggested that observers learn the potentially arbitrary ways that stimulation changes with eye position, so that a given color actually corresponds to the learned patterns of sensorimotor contingencies. As a thought experiment, they suggested that if observers were repeatedly exposed to a light that was red when viewed directly but switched to green whenever they looked away, then they would eventually learn this contingency and experience the stimulus as a single, stable color. The present perspective is not incompatible with their general thesis, but has a different emphasis. By the present account, the physiological differences across the visual field are instead discounted by differences in adaptation that normalize each retinal location to the *same* external stimulus. If we did constantly expose the fovea and periphery to *different* external stimuli, then the two retinal regions would each normalize for a different distribution of stimuli, and thus lead to divergent visual experiences.

3. PATTERN-SELECTIVE ADAPTATION

The foregoing example might seem trivial were it not for the fact that adaptation adjusts our vision not only for the average lighting in scenes but also to the patterns of light. These adjustments are known as contrast adaptation (because they are driven by the variations in light rather than the mean light level) and as pattern-selective adaptation (because the sensitivity changes are specific to the spatial, temporal, and chromatic properties of the adapting pattern)⁴. Well-known examples of pattern-selective adaptation in form perception include the tilt⁸, spatial-frequency⁹, and figural¹⁰ aftereffects. For example, after viewing a line that is tilted clockwise, a vertical line appears tilted counterclockwise. In the temporal domain, a classic example is the motion aftereffect or waterfall illusion¹¹. After watching a fall flowing downward, the static scenery to the side appears to drift upward. For these stimuli the sensitivity changes arise in the visual cortex, the first level at which the visual system appears to analyze patterns¹². For example, in the primate visual pathway, neurons in striate cortex are the first to show selectivity for features like orientation and direction of motion¹³. Thus the striate is the earliest likely site for an orientation or motion-selective aftereffect. However, we can again understand these aftereffects as a consequence of adaptive adjustments in a set of mechanisms that serve to renormalize their responses for the currently viewed stimulus. For example, a clockwise-tilted line will stimulate the mechanisms tuned to clockwise orientations more, and thus "light adapt" these mechanisms so that their responses are attenuated. This induces a counterclockwise tilt or negative afterimage in a vertical pattern, because the distribution of responses encoding vertical is now skewed by the sensitivity loss within the adapted mechanisms.

The perceptual aftereffects of pattern adaptation are so striking and so easy to induce that they are one of the most popular tools in the study of perception. A great deal of what we think we know about visual coding has come from experiments measuring how that coding is altered by adaptation (even though we still do not know a great deal about the actual basis of the sensitivity changes.) For example, adaptation has played a central role in the development of the channel theory of vision, and has been used extensively to measure the number and stimulus selectivities of the cortical channels representing form, motion, and color¹². However, unlike light adaptation, the role that pattern adaptation may play in shaping our perception remains largely unexplored. The visual channels that have been identified by adaptation have typically come from using relatively simple visual stimuli (e.g. drifting or tilted or colored gratings) and relatively simple visual judgments. Much less is known about the "channels" encoding the more holistic and ecological properties of the world – the meaningful objects and attributes (as opposed to the "stuff" of early vision¹⁴) - that correspond to our subjective experience of seeing. If pattern adaptation has a role in understanding these higher-level aspects of vision, then we should see its influence on the kinds of natural and biologically relevant perceptual judgments that we routinely make as we look about us. The following sections explore this by considering how adaptation influences two important perceptual decisions that we are making all the time.

4. ADAPTATION AND THE PERCEPTION OF FACES

The ability to identify and remember an individual face is perhaps the pinnacle of perceptual capacities. Face recognition depends on interpreting physically subtle configural cues in images that vary drastically because of changes in viewpoint and lighting and in the facial surface itself (e.g. because of changes in expression). Processing these cues is thought to depend on high-level visual mechanisms in extrastriate regions specialized for face coding^{15,16}.

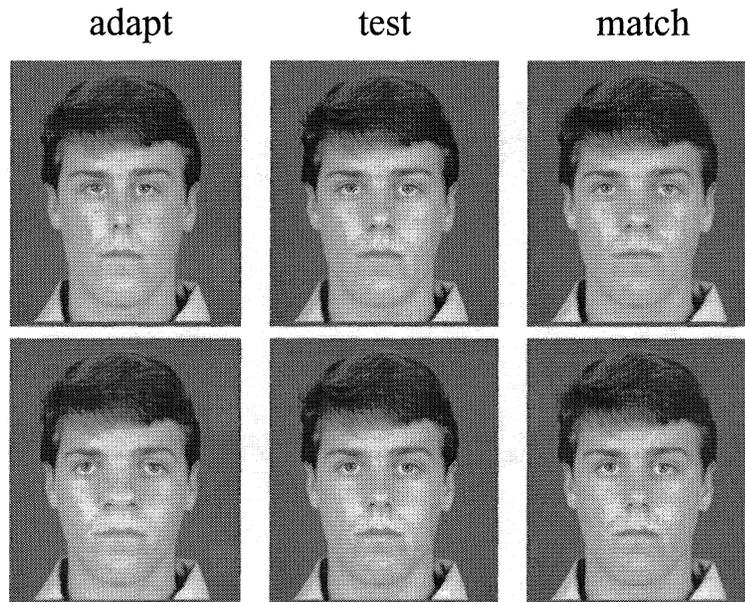


Figure 2. Face Adaptation. After viewing a contracted face (top left) the original face (middle) looks too expanded (top right). Adapting to an expanded face induces the opposite aftereffect (bottom).

We have conducted a series of experiments examining how the perception of faces is influenced by adaptation. Initially this was done by distorting an image of a face to create a new configuration, and then measuring how adaptation to the distortion would bias the appearance of the original image¹⁷. To quantify the bias, we created a large array of images that varied in the magnitude of the distortion. Observers then selected from this array the image that looked most like the original, before or after adaptation. The configural aftereffects we found were surprisingly strong, as illustrated by the example in Figure 2. Adaptation to a face that was locally contracted to form a pinched configuration caused the original face to look too expanded, and thus different from the remembered face. Consequently, the image that looked most like the learned face was now physically contracted. (That is, a physically contracted stimulus was required to null the perceptual expansion induced by the adaptation.) An expanded adapting image had the opposite effect.

Like other forms of adaptation, the perceptual distortions in faces are consistent with a process that renormalizes perception so that the adapting stimulus looks more neutral – or importantly, corresponds more closely to our expectations about the world. In the case of faces, this means that the physically distorted adapting face begins to appear more normal, and thus begins to define what a normal face is. In fact, we subsequently explicitly measured whether the face images in the array appeared normal (i.e. possible images of a real face) or distorted (i.e. not a possible face). These measurements showed that adaptation could strongly bias the gamut of images within the array that appeared acceptable to observers¹⁸.

Several lines of evidence suggest that the face adaptation alters sensitivity at a high and possibly face-specific level of processing. First, we found that the aftereffects were asymmetrical¹⁷. Distorted images biased the appearance of the original face, but not vice versa. In terms of low-level features, there is nothing special about the original image, and thus no reason why it should not be an equally effective adapter. Yet at the configural level it is special, because it conforms better to the distribution of faces we have been exposed to and to which we have been normalized. The distorted faces are instead like the red light source of Figure 1, and require readjustment. Second, two recent studies have shown that the aftereffects transfer across large changes in image size^{19,20}. And unlike the light adaptation effects we considered above, they also transfer across retinal location¹⁹. Thus they must in large part adjust to a property that is invariant to changes in position or scale, and this property is the facial configuration.

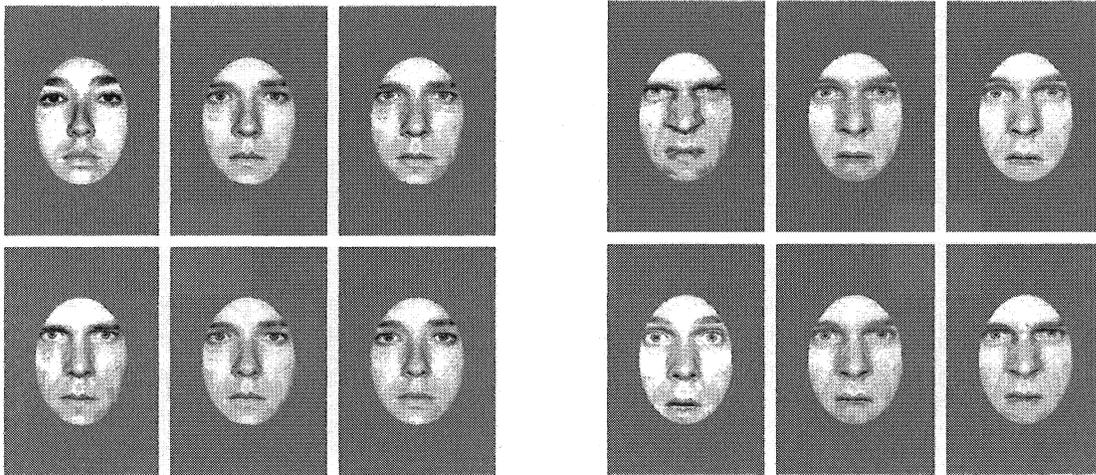


Figure 3. Adaptation-induced biases in gender or expression. Left: Adapting to a female face (top left) makes intermediate faces in a female-male morph appear more “male.” The original neutral face defining the gender boundary (middle) therefore appears more like the male face (top right), indicating that the neutral boundary image is now physically more like the female face. Opposite shifts are induced by adapting to the male face in the morph. Right: Similar aftereffects bias categorical ratings of expression.

The stimuli we originally used to probe the adaptation were distorted in arbitrary ways, and in particular do not capture the actual ways that real faces vary. However, Leopold et al¹⁹ showed comparable effects with stimuli drawn from empirically defined configural differences within a population of observers, thus showing that the normal variations we encounter are sufficient to induce very different adaptation states. With Daniel Kaping and Aaron Bilson, I have begun to examine how adaptation to images of actual faces biases the actual perceptual judgments we normally make when we look at faces²¹. These include categorical judgments of identity, gender, expression, and race. To explore these, we morph between two original images, and then measure the point along the morph at which ratings change from one category to another. Once again the biases induced by prior adaptation are dramatic. For example, Figure 3 shows the neutral points separating a male or female identification or between expressions. Prior adaptation to either extreme rapidly renormalizes face perception, biasing the original category boundary toward the opposite image and thus shifting the new neutral point toward the adapting image.

One important implication of these aftereffects is that they might reveal properties of the perceptual channels directly involved in face coding. However, the judgments are so easily perturbed that it raises the concern that we can adapt to anything, and thus that there is no constraint on the structures we might reveal. That is, perhaps we can adjust selectively to any arbitrary stimulus, and adaptation therefore reflects only this general malleability rather than any underlying skeleton. In this sense, examples where we fail to find an adaptive adjustment become important. Jill Yamashita in my lab has carried out a number of experiments testing the color and spatial selectivity of face adaptation²². Aftereffects are strongly selective for contrast polarity – for example, opposite aftereffects can be simultaneously induced in an image and its photographic negative. Yet the selectivity for color is by comparison very weak (though not entirely absent). That is, we cannot induce strongly separate spatial aftereffects for a red or green face, nor could we see any evidence for a color aftereffect that is contingent on a facial distortion (even though such color and form contingent aftereffects are well established for simple patterns²³, as in the McCollough Effect²⁴). Thus there may in fact be limits to the kinds of adaptive adjustments the visual system can make, and this means that the ones we can see may say something important about the intrinsic representational structure.

In either case, a second important implication is that adaptation can – and probably routinely does – strongly shape our face perception. Faces are interesting in this regard precisely because they vary across individuals. (If they did not, then they would not be very useful for recognition!) 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.

But can we see signs of these adaptation effects in our characteristic experience with faces? I think we can, and that adaptation may in fact provide a novel perspective for interpreting many aspects of human face perception. For example, faces that are physically average tend to be rated as more attractive²⁵. Yet, again, 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. This could 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. (In this context, it is interesting to note that the person depicted in Freud's portrait of the Queen looks much more like the depiction in his self-portrait than it does to an actual image of the Queen, but perhaps that is just my perception.) Adaptation could also provide a simple sensory explanation of how judgments of attractiveness could change when an individual is exposed to a new visual environment, as the following description of the anthropologist Malinowski's experiences suggests:

"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 196²⁶

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.

A second 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 groups²⁷. One reason why it is important for adaptation to normalize visual responses is so that the dynamic range of visual mechanisms is appropriate for the range of stimuli that the channels are likely to encounter. Mismatches are costly because they reduce the information that can be carried by the channel. For a Gaussian-like distribution of inputs, the optimal response curve is sigmoidal, so that discrimination is finest around the average input value and asymptotes for extreme values²⁸. In color a perceptual consequence of this is that we are very good at discriminating small changes around the white point, while poor at discriminating colors far from white (e.g. when trying to tell apart two saturated reds)²⁹. Moreover, light adaptation provides color with a "roving" neutral point so that optimal discrimination can always be centered around the average color. This is relevant to face coding if the processes mediating face perception are similarly matched by adaptation to the gamut of faces encountered. In that case they would be optimized for coding the differences between faces around the average. This predicts that for physically equivalent stimulus differences, judgments of facial similarity would increase as one moved farther from the average face, a behavior consistent with the other-race effect. It also predicts that sensitivity could be repositioned around a new face (or race) simply by adapting the observer to the new average. Thus, the other-race effect could be accounted for by the limited dynamic range characteristic of visual coding. Time will hopefully tell whether accounts of this kind are correct or ludicrous. But the point remains that even such complex perceptual judgments as beauty or facial similarity could in principle depend on relatively simple and well-established principles of visual coding and how it is adjusted through adaptation.

With these thoughts in mind, we can now return to Freud's portrait. The very act of exploring the face as he brought it to life 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 were struck by salient distortions, in his eye the same configural properties might have been 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 distinctiveness over time. For example, the distortions we tested often induced laughter when they were first presented to subjects, yet by the end of the session they appeared mundane. This analysis is not an aesthetic judgment about the work, nor does it preclude the possibility that the artist chose to emphasize or perhaps exaggerate certain traits that are perceptually obvious to him. I use it merely as an example to say that, if we could look at the painting through his eyes, it would not look the same.

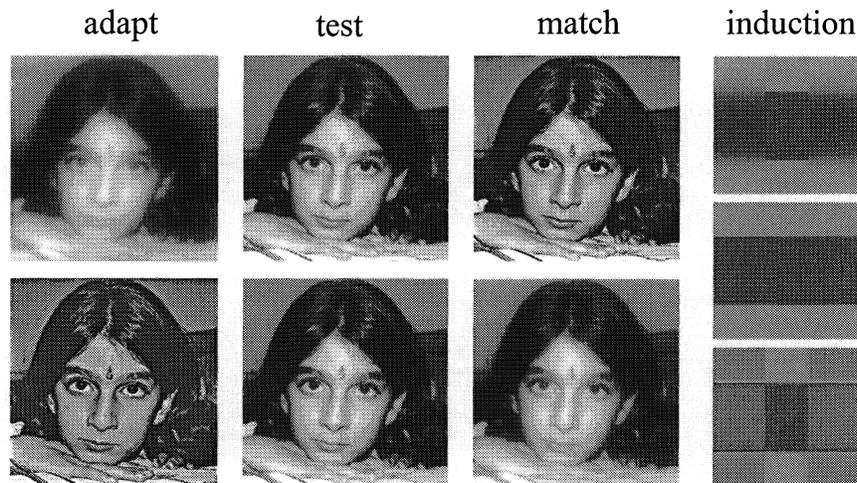


Figure 4. Adaptation to blur. Viewing a blurred image (top left) causes a physically focused image (middle) to appear too sharp (top right). A sharpened adapter leads to the opposite aftereffect. Similar effects are induced by placing an image in a blurred or sharpened surround. The edges in the center of each square are identical, but are blurred by a sharpened surround and sharpened by a blurred surround.

5. ADAPTATION AND THE PERCEPTION OF BLUR

The preceding section explored the possibility that individuals see the world differently because they are adapted to different worlds. This section considers whether physiologically different individuals may see in similar ways if they are adapted to the same world. This will be illustrated by examining how adaptation influences the perception of blur in images.

Blur is one of the most important dimensions of image quality, and like face recognition, is a stimulus that we make intuitive and natural judgments about all of the time. Unlike faces, there is no question that blur is also a property to which the visual system is constantly adjusting. Almost all studies examining these adjustments have focused on the accommodative changes in the eye's optics. Yet the neural visual system can also adjusting to the blur in the retinal image formed by the optics^{30,31}. We have recently begun to study these neural adjustments^{32,33}. The methods are very similar to the ones we used to examine adaptation to faces. Images were blurred or sharpened by changing the slope of their amplitude spectrum. Subjects first adapted by viewing these biased images for a few minutes, and then adjusted the spectral slope of target images until the target appeared properly focused.

The aftereffects in this case were once more surprisingly dramatic. After viewing a blurry image, a physically focused image appears too sharp, so that the point of subjective focus shifts toward a physically blurred image (Figure 4). Sharpened images induced the opposite aftereffects, causing the correctly focused image to appear blurred. Once again, these effects are consistent with a renormalization of our form perception, so that the currently viewed stimulus becomes the new prototype for proper image focus. In fact, this was again clear from observations on how the adapting stimuli themselves changed. Images that appeared strongly blurred or sharpened when we first presented them, often changed rapidly over the course of viewing so that they appeared better focused. While we have examined only short-term changes, there is also evidence for adaptive changes over much longer-term time scales. Fine, Smallman, and MacLeod³⁴ examined blur perception in an individual who had cataracts most of his life. After removal, edges appeared to him to be too sharp, and this "aftereffect" showed little sign of diminishing even after months. The time scales relevant for understanding visual adaptation are important because they tell us over what time scale we should be characterizing the stimulus distributions in the environment. However, this is an aspect of adaptation about which we still know little.

With regard to the themes I have been considering, there are two important aspects of the blur adaptation. The first is that we have better clues in the case of blur about the "prior" that the visual system is adjusting to (at least relative to the

priors involved in processes like face recognition). A number of studies have shown that the spatial statistics of natural images have a characteristic property. 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^{35,36}. 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 ways³⁶⁻³⁸. One reason this match is important is because it evenly distributes the responses across cortical cells tuned to different spatial scales³⁶. A convincing perceptual demonstration of this can be seen when we look at filtered noise patterns. As Field and Brady³⁹ 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.

The second important aspect is that blur is a case where we have a good understanding of how individuals differ, and where we can be confident that these physiological differences are very large. Specifically, we know that blur in the retinal image can vary widely because of differences in refractive errors. Yet individuals are often unaware of the errors in their optics, and when they become aware these are experienced as a failure of acuity rather than as a subjective experience that the world is out of focus⁴⁰. Arguably, this is because adaptation adjusts neural responses to maintain the balance of cortical responses across spatial scale, thus compensating for the imbalances in the image. Obviously, this cannot equate the visual experience of different observers, since the emmetropic observer will still be able to see details in the image that an incorrectly refracted observer cannot. More generally, adaptation cannot restore any information that is lost because of a difference between observers. Yet it can still bring them toward a more common experience.

It turns out that the processes underlying the changes in blur perception are again adjusting to higher-level aspects of the image. For example, the actual perception of blur does not depend simply on the overall amplitude spectrum. Images that are in focus can have different spectral slopes because they differ in the density of structure at different scales (e.g. the number of edges), and not in the amplitude of the structure (e.g. the blur in each edge)³⁹. And observers are very good at setting the correct focus for these images^{39,41}. Thus they are not simply equating the overall spectrum. Similarly, we have shown that how blur adaptation transfers across different images is better predicted by the actual focus of the images than by the differences in their amplitude spectra³³. This suggests that the visual system is not simply adapting independently to the contrasts present at each spatial scale. Rather, it is adjusting to some explicit representation of edge blur. But whatever its basis, the point here is that individuals with large and well defined differences in the properties of their visual systems may converge toward a common visual experience of the world (e.g. of blur), because they are adapted to common property of the world.

Like color, we can again 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 blurred⁴². Adaptation could readily compensate for the optical and physiological variations because the responses at each retinal region will be normalized for the same external stimulus. We in fact also observed a “simultaneous contrast” analog of the blur adaptation, which illustrates the active adjustments our vision is making to blur at different retinal locations⁴³. In Figure 4, the central bars are all identical squarewave edges. Yet the bars surrounded by blurry edges appear sharpened, while the bars surrounded by sharpened edges appear blurred. (Similar effects are induced when the bars are replaced by complex or natural images, and thus the effects do not depend on the alignment of local edges⁴³. It is also notable that we did not observe a similar spatial interaction affecting facial configuration, perhaps in part because the ability to accurately judge these configurations is restricted to the fovea.) How should we adjust the surrounding bars in Figure 4 so that the central bar appears properly focused? The answer is not completely obvious, because if we choose a physically focused surround then the “effective” surround might be blurred in the peripheral representation. Yet the contrast effects are nulled with the physically correct surround, showing that the interactions are already adjusted for the physiological variations with eccentricity⁴³.

6. ADAPTATION AND THE CONTENTS OF CONSCIOUSNESS

The processes of adaptation I have considered do not resolve the problem of perceptual qualia. For example, they do not provide an answer to the “inverted spectrum” problem – whether what I experience as red would look green to me through another’s eyes³. But they do place important constraints on color experience by defining the stimulus that looks gray, and they should similarly affect many if not all other perceptual experiences by defining the “grays” for those perceptions.

Adaptation may also fundamentally influence our conscious experience in further way. Barlow⁴⁴ has argued that adaptation is best viewed as 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 highlight the new associations, or “suspicious coincidences” in the world⁴⁵. 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 the world is always holding us under its spell. We are *always* adapted to it, and thus we are *always* experiencing an adaptation effect. Thus much of the content that reaches our conscious awareness may be a visual aftereffect.

7. 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 paper I have argued that this private experience is shaped in important ways by processes of adaptation, and I 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). I have also argued that we are always adapted to the 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, 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 coding⁴⁶. Ultimately these measurements may also reveal much about visual experience.

ACKNOWLEDGMENTS

I am grateful to Daniel Kaping for providing the images shown in Figure 3. Supported by EY-10834.

REFERENCES

1. S.M. Anstis, "Was El Greco astigmatic?," *Investigative Ophthalmology and Visual Science* 37, pp. S697, 1996.
2. C.C. Moore, A.K. Romney, T. Hsia, "Shared cognitive representations of perceptual and semantic structures of basic colors in Chinese and English," *Proceedings of the National Academy of Science* 97, pp. 5007-5010, 1997.
3. S.E. Palmer, "Color, consciousness, and the isomorphism constraint," *Behavioral and Brain Sciences* 22, pp. 923-989, 1999.
4. M.A. Webster, Pattern-selective aftereffects in color and form perception, in: L. Chalupa, J.S. Werner (Eds.), *The Visual Neurosciences* (MIT Press, Cambridge, forthcoming).
5. M.A. Webster, "Human colour perception and its adaptation," *Network: Computation in Neural Systems* 7, pp. 587-634, 1996.
6. I. Abramov, J. Gordon, H. Chan, "Color appearance in the peripheral retina: Effects of stimulus size," *Journal of the Optical Society of America A* 8, pp. 404-414, 1991.
7. J.K. O'Regan, A. Noe, "A sensorimotor account of vision and visual consciousness," *Behavioral and Brain Sciences*, in press.
8. J.J. Gibson, M. Radner, "Adaptation, after-effect and contrast in the perception of tilted lines. I. Quantitative studies," *Journal of Experimental Psychology* 20, pp. 453-467, 1937.
9. C. Blakemore, P. Sutton, "Size adaptation: A new aftereffect," *Science* 166, pp. 245-247, 1969.
10. W. Kohler, H. Wallach, "Figural aftereffects: An investigation of visual processes," *Proceedings of the American Philosophical Society* 88, pp. 269-357, 1944.
11. G. Mather, F. Verstraten, S. Anstis (Eds.), *The Motion Aftereffect*, MIT Press, Cambridge MA, 1998.
12. N. Graham, *Visual Pattern Analyzers* Oxford University Press, Oxford, 1989.
13. R.L. De Valois, K.K. De Valois, *Spatial Vision* Oxford University Press, Oxford, 1988.

14. E.H. Adelson, J.R. Bergen, The plenoptic function and the elements of early vision, in: M.S. Landy, J.A. Movshon (Eds.), *Computational Models of Visual Processing* (MIT Press, Cambridge, 1991) 3-20.
15. M.J. Farah, K.D. Wilson, M. Drain, J.N. Tanaka, "What is "special" about face perception?," *Psychological Review* 105, pp. 482-498, 1998.
16. N. Kanwisher, "Domain specificity in face perception," *Nature Neuroscience* 3, pp. 759-763, 2000.
17. M.A. Webster, O.H. MacLin, "Figural after-effects in the perception of faces," *Psychonomic Bulletin and Review* 6, pp. 647-653, 1999.
18. O.H. MacLin, M.A. Webster, "Influence of adaptation on the perception of distortions in natural images," *Journal of Electronic Imaging* 10, pp. 100-109, 2001.
19. D.A. Leopold, A.J. O'Toole, T. Vetter, V. Blanz, "Prototype-referenced shape encoding revealed by high-level aftereffects," *Nature Neuroscience* 4, pp. 89-94, 2001.
20. L. Zhao, C.F. Chubb, "The size-tuning of the face-distortion aftereffect," *Vision Research* 41, pp. 2979-2994, 2001.
21. D. Kaping, A. Bilson, M.A. Webster, "Adaptation and Categorical Judgments of Faces," 2002 Vision Sciences Society Abstract, submitted.
22. J.A. Yamashita, J.L. Hardy, K.K. De Valois, M.A. Webster, "Color and contrast-selective after-effects in the perception of faces," *Investigative Ophthalmology and Visual Science* 41, pp. S224, 2000.
23. C.F.I. Stromeyer, Form-color aftereffects in human vision, in: R. Held, H.W. Leibowitz, H.L. Teuber (Eds.), *Handbook of Sensory Physiology*, Vol. VIII (Springer-Verlag, New York, 1978).
24. C. McCollough, "Color adaptation of edge-detectors in the human visual system," *Science* 149, pp. 1115-1116, 1965.
25. J.H. Langlois, L.A. Roggman, "Attractive faces are only average," *Psychological Science* 1, pp. 115-121, 1990.
26. D. Symons, *The Evolution of Human Sexuality* Oxford University Press, New York, 1979.
27. T. Valentine, "A unified account of the effects of distinctiveness, inversion, and race in face recognition," *Quarterly Journal of Experimental Psychology* 43A, pp. 161-204, 1991.
28. S.B. Laughlin, "Form and function in retinal processing," *Trends in Neuroscience* 10, pp. 478-483, 1987.
29. T. von der Twer, D.I.A. MacLeod, "Optimal nonlinear codes for the perception of natural colours," *Network: Computation in Neural Systems* 12, pp. 395-407, 2001.
30. M. Mon-Williams, J.R. Tresilian, N.C. Strang, P. Kochhar, J.P. Wann, "Improving vision: neural compensation for optical defocus," *Proceedings of the Royal Society of London B* 265, pp. 71-77, 1998.
31. M.A. Webster, E. Miyahara, "Contrast adaptation and the spatial structure of natural images," *Journal of the Optical Society of America A* 14, pp. 2355-2366, 1997.
32. M.A. Webster, Contrast sensitivity under natural states of adaptation, *SPIE Proceedings* 3644, pp. 58-70, 1999.
33. M.A. Webster, S.M. Webster, J. MacDonald, S.R. Bharadwaj, "Adaptation to blur," *SPIE Proceedings* 4299, pp. 69-78, 2001.
34. I. Fine, H.S. Smallman, D.I.A. MacLeod, "Visual processing before and after removal of bilateral congenital cataracts in adulthood: contrast constancy and edge perception," *Investigative Ophthalmology and Visual Science* 41, pp. S995, 2000.
35. G.J. Burton, I.R. Moorhead, "Color and spatial structure in natural scenes," *Applied Optics* 26, pp. 157-170, 1987.
36. D.J. Field, "Relations between the statistics of natural images and the response properties of cortical cells," *Journal of the Optical Society of America A* 4, pp. 2379-2394, 1987.
37. J.J. Atick, A.N. Redlich, "What does the retina know about natural scenes?," *Neural Computation* 4, pp. 196-210, 1992.
38. M.V. Srinivasan, S.B. Laughlin, A. Dubs, "Predictive coding: A fresh view of inhibition in the retina," *Proceedings of the Royal Society of London B* 216, pp. 427-459, 1982.
39. D. Field, N. Brady, "Visual sensitivity, blur, and the sources of variability in the amplitude spectra of natural images," *Vision Research* 37, pp. 3367-3383, 1997.
40. R.J. Watt, "An outline of the primal sketch in human vision," *Pattern Recognition Letters* 5, pp. 139-150, 1987.
41. Y. Tadmor, D.J. Tolhurst, "Discrimination of changes in the second-order statistics of natural and synthetic images," *Vision Research* 34, pp. 541-554, 1994.
42. S.J. Galvin, R.P. O'Shea, A.M. Squire, D.G. Govan, "Sharpness overconstancy in peripheral vision," *Vision Research* 37, pp. 2035-2039, 1997.
43. S.M. Webster, M.A. Webster, J. Taylor, J. Jaikumar, R. Verma, "Simultaneous blur contrast," *SPIE Proceedings* 4299, pp. 414-422, 2001.

44. H.B. Barlow, A theory about the functional role and synaptic mechanism of visual after-effects, in: C. Blakemore (Ed.), *Vision: Coding and Efficiency* (Cambridge University Press, Cambridge, 1990) 363-375.
45. H.B. Barlow, "Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception," *Vision Research* 30, pp. 1561-1571, 1990.
46. E.P. Simoncelli, B.A. Olshausen, "Natural image statistics and neural representation," *Annual Review of Neuroscience* 24, pp. 1193-1216, 2001.