Optimizing visual performance by adapting images to observers

Michael A. Webster*[a], Igor Juricevic[b]

[a]Dept. of Psychology/296, University of Nevada Reno, Reno, NV 89557, USA; [b]Dept. of Psychology, Indiana University, South Bend, IN 46634, USA

ABSTRACT

Visual adaptation is widely assumed to optimize visual performance, but demonstrations of functional benefits beyond the case of light adaptation remain elusive. The failure to find marked improvements in visual discriminations with contrast or pattern adaptation may occur because these become manifest only over timescales that are too long to probe by briefly adapting observers. We explored the potential consequences of color contrast adaptation by instead "adapting" images to simulate how they should appear to observers under theoretically complete adaptation to different environments, and then used a visual search task to measure the ability to detect colors within the adapted images. Color salience can be markedly improved for extreme environments to which the observer is not routinely exposed, and may also be enhanced even among naturally occurring outdoor environments. The changes in performance provides a measure of how much in theory the visual system can be optimized for a given task and environment, and can reveal the extent to which differences in the statistics of the environment or the sensitivity of the observer are important in driving the states of adaptation. Adapting the images also provides a potential practical tool for optimizing performance in novel visual contexts, by rendering image information in a format that the visual system is already calibrated for.

Keywords: Adaptation, visual salience, color vision, visual search, image processing, natural images

1. INTRODUCTION

A central premise of vision science is that the visual coding is adapted to match the visual properties of the organism's environment[1]. These adaptations occur over multiple timescales, from phylogeny to ontogeny to the transient adjustments when the lighting changes. Much of psychophysics has focused on brief adaptation timescales by exposing observers to a stimulus and then measuring how sensitivity and perception changes. The wealth of these aftereffects reveals that visual coding is highly malleable and thus may be calibrated for many attributes of the world[2]. A further central assumption is that these calibrations optimize perception and performance[3-5]. Yet the full functional consequences and benefits of what is conventionally considered as perceptual adaptation remain poorly characterized. In particular, it remains unclear to what extent and in what ways this exposure-dependent adaptation matches an individual's perception to their visual diet, and what performance potential this match affords them.

Empirically answering these questions is difficult in part because the most important consequences of adaptation may only become manifest over durations that are too long in practice to actually measure in observers. There is in fact growing evidence for longer and potentially distinct forms of adaptation from the rapid aftereffects that are conventionally studied[6-10]. However even these experiments are restricted to hours or days of exposure, and thus may fail to reveal the full impact of the adjustments. One solution to this problem would be to characterize variations in the statistical structure of the world and correlate this with variations in visual sensitivity. A potential example of this is studies of the "other-race" effect in face perception, where individuals are typically better at discriminating members of their own social group than other ethnicities, plausibly because their face perception is adapted to the specific distribution of faces they encounter[11, 12]. However, in such cases it is impossible to know what the actual stimulus history has been, and to know whether the effects reflect processes of visual adaptation or other factors such as learning.

We recently developed a novel approach to address the consequences of long-term adaptation to the visual environment, by "adapting" the image rather than the observer[13]. To the extent that basic mechanisms of visual coding are known and that reasonable inferences can be made about how the responses of these mechanisms are adapted, an image can be adjusted to simulate how it should appear to an observer adapted to any given stimulus context. This allows us to simulate theoretically complete states of adaptation that would be very difficult to probe by adapting observers. Evaluating perception and performance with these images can then be used to predict how much variation in visual coding can be attributed to variations in the environment, and how important it is for different perceptual tasks to be correctly calibrated for the ambient environment.

Obviously, this approach poses its own problems from the outset, because the nature of visual mechanisms and of the response changes with adaptation are themselves very poorly characterized[2], as are the relevant statistics of the environment[14]. However, we have begun by applying the approach to the specific context of color perception. Compared to other stimulus dimensions, the early stages of color coding are fairly well established, as are the patterns of adaptation to different distributions of colors[15]. Moreover, there are a growing number of databases of natural color images that have been used to characterize different environments[16-20]. Thus in the case of color reasonable assumptions can be made about what kinds of color environments observers might be immersed in and about how adaptation to these environments might adjust color coding. Importantly, color across environments significantly varies – both in different places and at different times (e.g. when the seasons change) – and this predicts that observers experience long-term exposure to a variety of color worlds[16, 20].

In previous studies, we used models of color adaptation to different naturalistic environments to explore how much color appearance might vary because of how much the world itself varies. This allowed us to characterize to what extent perception might change when the same observer is placed in different environments[21], or different observers (with different spectral sensitivities) are placed in the same environment[22]. In the present study, we instead examine the long-term consequences of adaptation on visual performance, by measuring how the salience of colors in an environment is affected by theoretically complete adaptation to that environment.

2. METHODS

2.1 Model visual system. The model we used for simulating color adaptation has been described in detail previously[13, 21, 22], and is based on a simple yet standard model of human color processing (Figure 1). Light is absorbed in the three classes of cones (L, M, and S), whose signals are then combined with different weights to form different postreceptoral channels. The channels shown correspond to the three "cardinal" mechanisms in the retina and LGN, which are sensitive to LvsM, SvsLM or luminance contrast[23]. However, these mechanisms appear elaborated to form multiple "higher-order" mechanisms in the cortex, so that color is instead represented by a large number of channels with varying spectral sensitivities[24]. Part of the evidence for this comes from studies of contrast adaptation, which can be selective for any arbitrary combination of luminance and chromatic contrast and therefore cannot be explained by independent response changes in only three channels[25, 26]. To incorporate this property we sampled the volume of color space with 26 mechanisms, spaced at 45 deg intervals in a scaled version of cardinal axis space[21]. The color of the stimulus was rendered by taking the vector sum of the channel outputs along the 3 cardinal directions, with the responses normalized so that these summed outputs returned the original RGB pixel values for the reference environment.
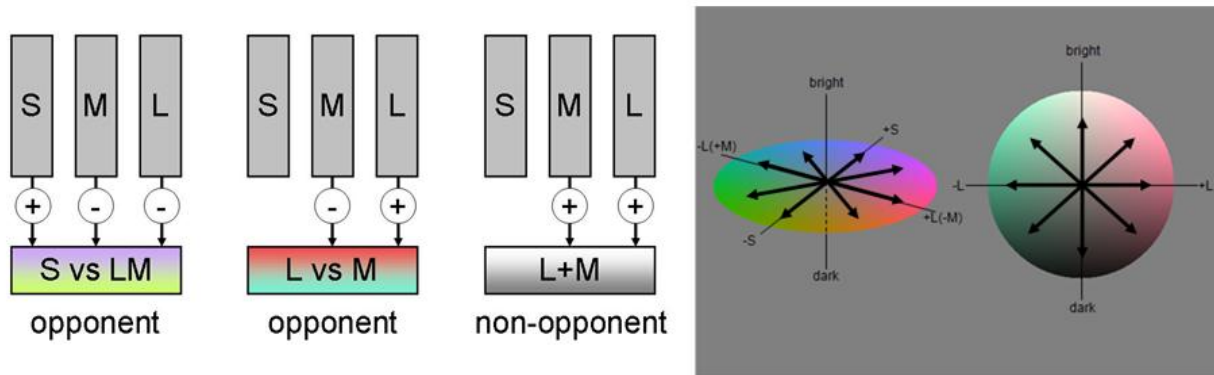

Figure 1. Image color was encoded by short, medium, and long-wave sensitive cone receptors (S,M,L) and then combined to form opponent chromatic channels or non-opponent luminance channels. The full set of post-receptoral channels included 26 mechanisms with preferred color directions at intervals of 45 degs within the luminance and chromatic planes .Colors were rendered as the vector sum of the adapted channel outputs along SvsLM, LvsM, and achromatic axes.

2.2 Adaptation. Adaptation was assumed to independently adjust the gain of each channel so that the average response to the current environment was the same as to the reference environment. At the level of the cones this adjusts to the average color in the scenes, while at the level of the post-receptoral channels it instead adjusts to the range of colors or to the contrast. Independent multiplicative gain changes are well described in the cones (von Kries adaptation)[27-30]. At short timescales the effects of contrast adaptation instead appear to reflect changes in contrast gain and result in a roughly
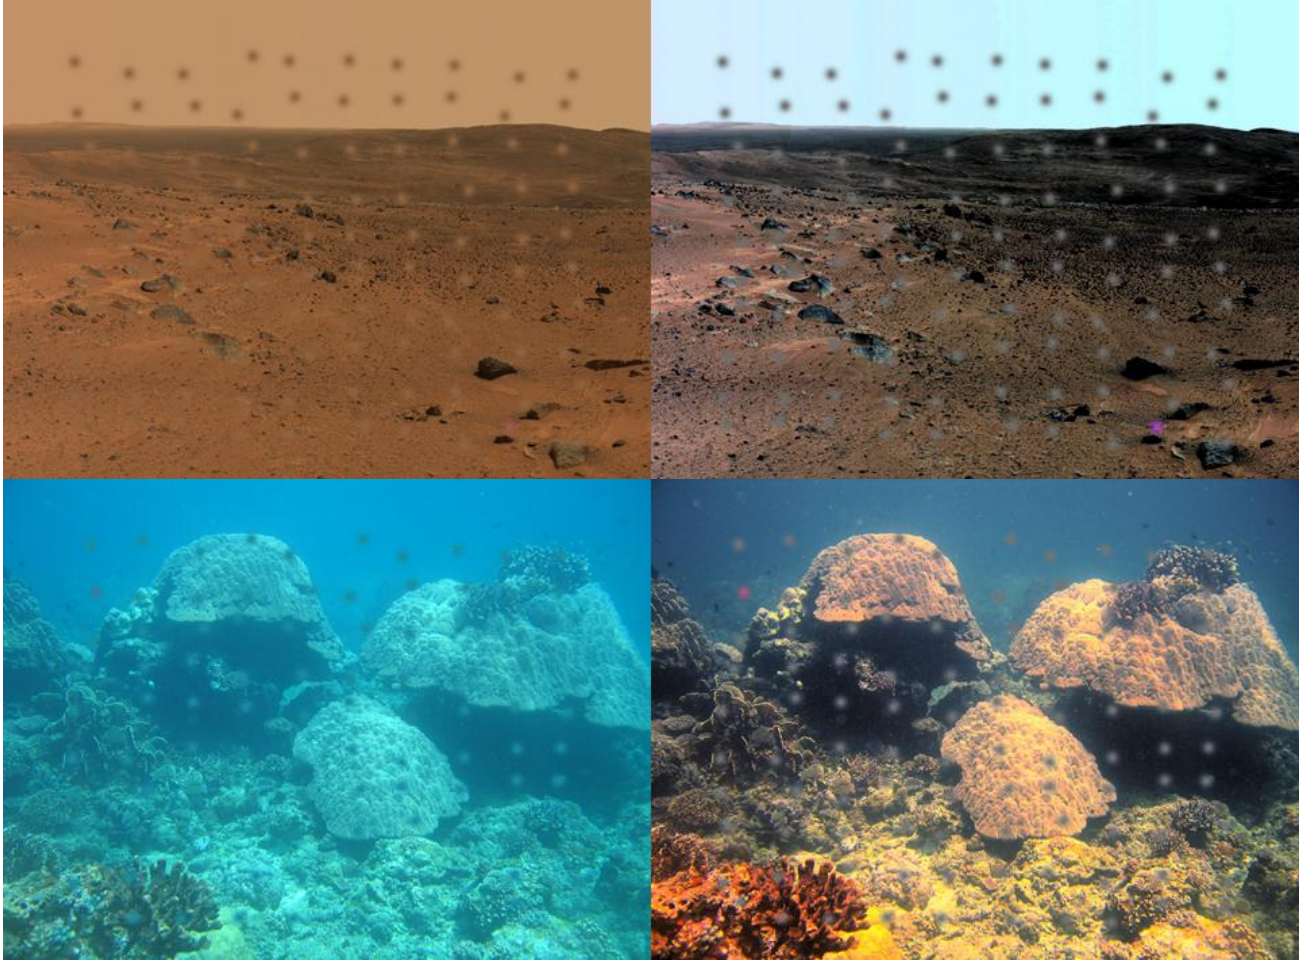
Figure 2. Examples of the original and adapted images from the scenes of the Martian surface or underwater environment. The left images are as seen by an observer adapted to land on earth, while the right images depict what the scenes would look like to an observer adapted to each depicted environment. Added spots show the set of distractors and single target color which appeared at a random location. Both the search stimuli and the background images differ only in the simulated state of adaptation. The target colors (e.g. purplish spot in the lower right corner of the Mars scene; red spot in the upper left corner of the aquatic scene) are conspicuous in the adapted images while difficult to resolve in the original unadapted images.

subtractive change in perceived contrast[26, 31]. However, at longer timescales there is recent evidence that the sensitivity changes may instead involve changes in response gain[10]. The latter is consistent with the assumption that each channel should match its dynamic range to the range of levels in the stimulus[13]. This is also consistent with the relative contrast sensitivity for luminance and chromatic contrast, which is much higher for chromatic mechanisms because the range of available cone contrasts is much less[32]. Accordingly, we assumed each post-receptoral channel also adapts independently by multiplicative gain changes.

2.3 Stimuli. Images included calibrated image sets of natural outdoor scenes[16] as well as uncalibrated sets collected for more "extreme" environments from the internet. In this study we show results for two of these environments – underwater scenes and the surface of Mars. For each set we calculated the average response within each channel to all pixels within all of the sampled images. The pixel values in the adapted images were then rendered after rescaling this average response to match the average response of the channel to the reference set.

For the visual search task, we added an irregular 10 by 10 array of spots onto the original images (Figure 2). The color of each spot was a weighted average of a single chromaticity and the underlying background, with the weight varying as a Gaussian (so that the chromaticity faded into the background with increasing distance from the center). For distractors, the luminance and chromaticity were all set to the average of the image. Targets were instead varied over a range of

colors relative to this average. Colors for both the images and the added array were then adjusted with the adaptation algorithm, so that there were two pairs of stimuli that differed only in the simulated state of adaptation. In practice, this was done by first adding the targets to the adapted image so that they corresponded to a fixed and known set of color angles and contrasts (spaced at 45 deg intervals in the LvsM and SvsLM chromatic plane), and then adjusting back to the original image (within which the targets had variable contrasts and hue angles).

2.4 Procedure. Observers searched for the targets within both the original images and the adapted images shown in random interleaved order. A button box was used to indicate the target location (quadrant), with the reaction time recorded for correct responses or timed out at 5 sec if the target could not be detected. Results reported are based on the average reaction times for multiple image pairs (adapted vs. original) for a given environment. Observers included author IJ and undergraduate students who participated for course credit. All had normal color vision as assessed by standard screening tests.
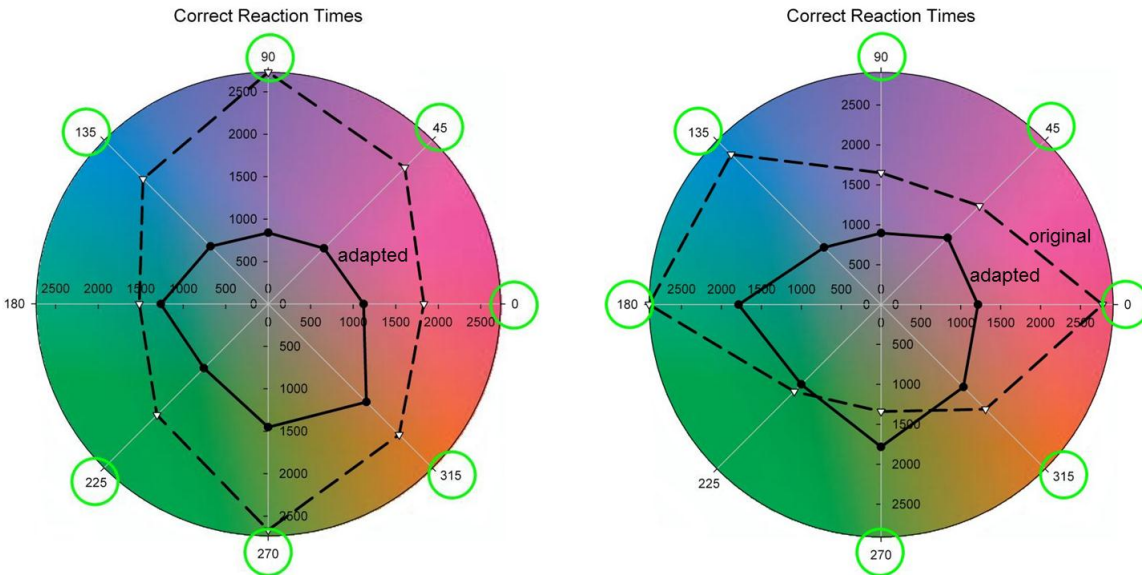

Figure 3. Mean reaction times for detecting the target colors in the adapted (solid line) or original (dashed line) images of Mars (left) or underwater (right). Times are plotted as the distance from the plot center, for target hues indicated by the direction from the center. Green circles along the circumference indicate target color directions for which search times were significantly faster in the adapted images.

3. RESULTS

Figure 3 compares average reaction times for detecting targets along different color directions either in the original images or the adapted images for the Mars (left) or underwater (right) environments. Note again that for each environment the image pairs were identical except for the simulated state of adaptation, and that the performance changes thus reflect the improvements expected if observers are under theoretically complete adaptation to the environment, given the specific assumptions we made about the adaptation. For both of these extreme contexts the improvements are obvious, and in fact underestimate the benefits of the adaptation, for the reported averages include many trials with the original images where the target could not be correctly identified within the 5 sec trial limit. Note also that the color directions that show increased salience differ for the two environments. For the Mars images the sensitivity gains are strongest along the SvsLM axis, while for the aquatic scenes the contrast reductions and thus response gains are stronger along the reddish-greenish dimension. In both cases the increased salience is consistent with an expanded gamut of color in the adapted images. For example, Figure 4 shows that the ratio of reaction times decreased roughly in proportion to the increase in the contrast ratio of the target under the two adapted states.

Figure 5 shows comparable settings when switching between two natural terrestrial environments[16]. In this case the original images were drawn from outdoor scenes in India during the monsoon season and thus from a lush green environment. These were then adapted to the distribution of colors sampled for outdoor vicinity of the more arid Reno, Nevada environment of our participants. Color contrasts in the images and thus search times are altered less in this case,

yet there were still significant improvements along some color directions. This suggests that even among the natural environments that we might routinely encounter, there is sufficient variation in the color distributions that adaptation can impact visual performance.
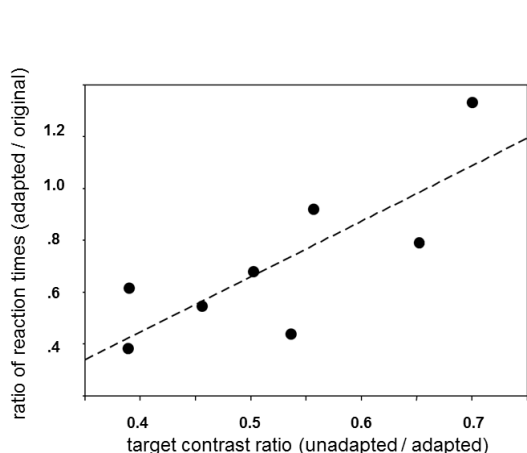

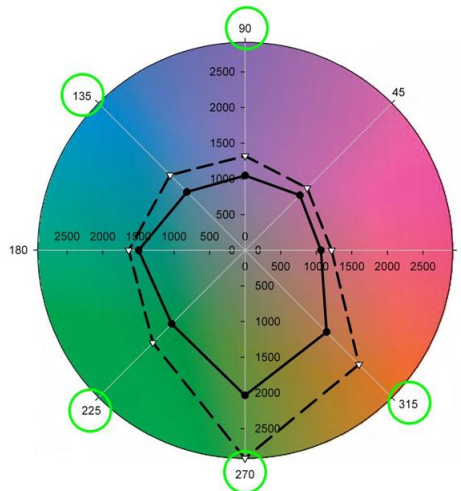Figure 4. Change in search times vs. change in contrast

Figure 5. Mean reaction times for detecting targets in lush scenes before (dashed) or after (solid) adaptation to arid scenes. Symbols as in Figure 3.

4. DISCUSSION

The finding that expanding the color gamut in low contrast images makes it easier to detect a color difference is not surprising, but is also not trivial, for it again reflects the putative functional role of adaptation in matching visual responses to the visual world, and the consequences this has for stimulus salience. Moreover, the adaptations to the extreme environments we tested are arguably no less extreme than the initial calibrations the visual system must make during development. Individual differences in spectral sensitivity are enormous and arise from many factors, from the density of preretinal screening pigments[33] to the pronounced variation in the ratio of the different cone classes[34] to less well quantified variations in the number and nature of post-receptoral mechanisms. Yet these factors have very little impact on measures of color appearance, demonstrating that to a large extent appearance is compensated for idiosyncratic characteristics of the observer[35-37]. Moreover, even within the individual there are profound changes in the optical and neural properties of the visual system throughout the lifespan. Yet again color appearance remains largely impervious to the insults of age, revealing that the visual system undergoes a continuous recalibration to maintain perceptual constancy[38, 39]. (In fact, the processes of adaptation themselves appear to remain largely robust to aging, potentially because they are fundamentally important for maintaining visual function[40, 41].)

Again, the full extent of these adjustments is difficult to directly probe in practice because they may require very extended time and exposure to unfold. The approach we used provides a way to partially circumvent this problem, by allowing changes in visual performance to be tested under conditions in which the adaptation is in theory complete. This provides a measure of the potential limits of functional consequences of adaptation, but is itself limited by the model of adaptation. One such factor is that our model assumes that color salience and color appearance are equivalent – that is, the colors that are more noticeable are more saturated. However, the goals of detecting novel features in a search task and maintaining perceptual constancy are not necessarily compatible[42], and it is very likely that there are other facets of how adaptation alters the perception of colors that our model does not capture. Another limitation is that we assumed that adaptation altered the signal but not the noise. If the noise instead occurs prior to the gain change, then the adjustments might alter appearance much more than discriminability[43, 44].

The specific conditions we tested suggest there can be significant consequences of color contrast adaptation on visual performance even within the range of natural outdoor environments that humans are normally exposed to. Again this was hinted at by the finding that participants in an arid environment were better at searching for colors adapted to that environment compared to a lush environment. This is admittedly circumstantial, and more compelling evidence would

require testing the converse condition. However, even if performance does not differ across two contexts, we argue that this would represent an important clue to understanding why the visual system adapts. If the natural range of environmental variation is too small to require adaptation to optimize visual performance for a given goal, then this constrains functional accounts of the adaptation. Moreover, it also provides a way to explore the sources of variation that are actually important for the adaptation. The approach we developed could also be used to measure visual search within images that simulate observers with different spectral sensitivities and under different states of adaptation (e.g. comparing color search in young observer through their own eyes or in images simulating the brunescent lens of an old observer or vice versa). Variations in the physical properties of the world and the physiological properties of the observer are fundamentally different, but are nevertheless intimately linked through the processes of adaptation[2]. Analyzing their relative impact on perception and performance could thus help reveal which source of variation is more important for understanding the consequences of adaptation.

While we have emphasized the implications of this approach for testing functional theories of adaptation, the method of pre-adapting images also has many practical applications, for it allows the visual environment to be matched to the observer, obviating the extensive visual training that might normally be required for observers to adjust to new contexts. The surface of Mars and the floor of the sea are examples of worlds that we do increasingly explore, even if remotely, and ones which could in principle be visually transformed so that their properties are more like the worlds we are "used to" in order to optimize this exploration.

As noted, we adapted the color characteristics of the images because the mechanisms of color coding are relatively well defined. However, the principles can be extended to other dimensions such as the spatial structure of images, which could again be pre-adapted to match the structure observers are typically adapted to. For example, Kompaniez et al. [45] recently examined adaptation to mammograms, and whether prior adaptation can facilitate visual search for anomalous targets (e.g. tumors) in the images[46]. Medical images have characteristic texture patterns and are substantially blurrier than typical natural images, and both perception and performance were affected even by relatively brief periods of adaptation. Thus part of the "learning" that trained radiologists acquire may reflect extended adaptation to this background structure. Altering the images so that they are instead more like the spatial diet we normally encounter could provide a novel way of enhancing the salience of outliers in medical images, by discounting the salience of the background. This raises the possibility that performance in many visual contexts could be rapidly and optimally enhanced by applying corrections based on realistic models of visual adaptation.

5. ACKNOWLEDGMENTS

REFERENCES

[1] Atick, J.J., Could information-theory provide an ecological theory of sensory processing. Network - Computation in Neural Systems 3, 213-251 (1990).
[2] Webster, M.A., Adaptation and visual coding. J Vision, 11(5), 3: 1-23 (2011).
[3] Clifford, C.W., et al., Visual adaptation: neural, psychological and computational aspects. Vision Research 47, 3125-31 (2007).
[4] Kohn, A., Visual adaptation: physiology, mechanisms, and functional benefits. J Neurophysiol 97, 3155-64 (2007).
[5] Wark, B., B.N. Lundstrom, and A. Fairhall, Sensory adaptation. Curr Opin Neurobiol 17, 423-9 (2007).
[6] Kording, K.P., J.B. Tenenbaum, and R. Shadmehr, The dynamics of memory as a consequence of optimal adaptation to a changing body. Nat Neurosci 10, 779-86 (2007).
[7] Neitz, J., et al., Color perception is mediated by a plastic neural mechanism that is adjustable in adults. Neuron 35, 783-92 (2002).
[8] Zhang, P., et al., Effects of orientation-specific visual deprivation induced with altered reality. Curr Biol 19, 1956-60 (2009).
[9] Bao, M. and S.A. Engel, Distinct mechanism for long-term contrast adaptation. Proc Natl Acad Sci U S A 109, 5898-903 (2012).
[10] Kwon, M., et al., Adaptive changes in visual cortex following prolonged contrast reduction. J Vision 9, 20: 1-16 (2009).
[11] Armann, R., et al., Race-specific norms for coding face identity and a functional role for norms. J Vision, 11(13), 9 (2011).

[12] Webster, M.A. and D.I.A. MacLeod, Visual adaptation and face perception. Philos Trans R Soc Lond B Biol Sci 366, 1702-25 (2011).

[13] McDermott, K., et al., Adapting images to observers. SPIE 68060, V-1-10 (2008).

[14] Simoncelli, E.P. and B.A. Olshausen, Natural image statistics and neural representation. Annu Rev Neurosci 24, 1193-216 (2001).

[15] Webster, M.A., Human colour perception and its adaptation. Network: Computation in Neural Systems 7, 587-634 (1996).

[16] Webster, M.A., Y. Mizokami, and S.M. Webster, Seasonal variations in the color statistics of natural images. Network: Computation in Neural Systems 18, 213-33 (2007).

[17] Tkacik, G., et al., Natural images from the birthplace of the human eye. PLoS One 6, e20409 (2011).

[18] Nascimento, S.M., F.P. Ferreira, and D.H. Foster, Statistics of spatial cone-excitation ratios in natural scenes. J Opt Soc Am A 19, 1484-90 (2002).

[19] Ruderman, D.L., T.W. Cronin, and C.C. Chiao, Statistics of cone responses to natural images: implications for visual coding. J Opt Soc Am A 15, 2036-2045 (1998).

[20] Webster, M.A. and J.D. Mollon, Adaptation and the color statistics of natural images. Vision Research 37, 3283-3298 (1997).

[21] Juricevic, I. and M.A. Webster, Variations in normal color vision. V. Simulations of adaptation to natural color environments. Vis Neurosci 26, 133-45 (2009).

[22] Webster, M.A., I. Juricevic, and K.C. McDermott, Simulations of adaptation and color appearance in observers with varying spectral sensitivity. Ophthalmic Physiol Opt 30, 602-10 (2010).

[23] Derrington, A.M., J. Krauskopf, and P. Lennie, Chromatic mechanisms in lateral geniculate nucleus of macaque. Journal of Physiology 357, 241-65 (1984).

[24] Eskew, R.T., Jr., Higher order color mechanisms: a critical review. Vision Research 49, 2686-704 (2009).

[25] Krauskopf, J., et al., Higher order color mechanisms. Vision Research 26, 23-32 (1986).

[26] Webster, M.A. and J.D. Mollon, The influence of contrast adaptation on color appearance. Vision Research 34, 1993-2020 (1994).

[27] Brainard, D.H. and B.A. Wandell, Asymmetric color matching: how color appearance depends on the illuminant. J Opt Soc Am A 9, 1433-48 (1992).

[28] Chichilnisky, E.J. and B.A. Wandell, Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. Vision Research 35, 239-54 (1995).

[29] Wuerger, S.M., Color appearance changes resulting from iso-luminant chromatic adaptation. Vision Research 36, 3107-18 (1996).

[30] Webster, M.A. and J.A. Wilson, Interactions between chromatic adaptation and contrast adaptation in color appearance. Vision Research 40, 3801-16 (2000).

[31] Georgeson, M.A., The effect of spatial adaptation on perceived contrast. Spatial Vision 1, 103-12 (1985).

[32] McDermott, K.C. and M.A. Webster, The perceptual balance of color. J Opt Soc Am A 29, A108-17 (2012).

[33] Webster, M.A. and D.I. MacLeod, Factors underlying individual differences in the color matches of normal observers. J Opt Soc Am A 5, 1722-35 (1988).

[34] Hofer, H., et al., Organization of the human trichromatic cone mosaic. J Neurosci 25, 9669-79 (2005).

[35] Brainard, D.H., et al., Functional consequences of the relative numbers of L and M cones. J Opt Soc Am A 17, 607-14 (2000).

[36] Miyahara, E., et al., Color vision in two observers with highly biased LWS/MWS cone ratios. Vision Research 38, 601-12 (1998).

[37] Webster, M.A., et al., Variations in normal color vision. II. Unique hues. J Opt Soc Am A 17, 1545-55 (2000).

[38] Schefrin, B.E. and J.S. Werner, Loci of spectral unique hues throughout the life span. J Opt Soc Am A 7, 305-11 (1990).

[39] Werner, J.S. and B.E. Schefrin, Loci of achromatic points throughout the life span. J Opt Soc Am A 10, 1509-1516 (1993).

[40] Elliott, S.L., et al., Aging and blur adaptation. J Vision 7:6, 8 (2007).

[41] Elliott, S.L., J.S. Werner, and M.A. Webster, Individual and age-related variation in chromatic contrast adaptation. J Vision 12:8, 11 (2012).

[42] McDermott, K.C., et al., Adaptation and visual salience. J Vision 10:13, 17 (2010).

[43] Rieke, F. and M.E. Rudd, The challenges natural images pose for visual adaptation. Neuron 64, 605-16 (2009).

[44] MacLeod, D.I.A., Colour discrimination, colour constancy, and natural scene statistics (The Verriest Lecture), in Normal and Defective Colour Vision, J.D. Mollon, J. Pokorny, and K. Knoblauch, Editors. 2003, Oxford University Press: London.

[45] Kompaniez, E., et al., Adaptation and the perception of structure in radiological images (abstract). J Vision 11:1, 12 (2011).

[46] Kompaniez, E., et al., Adaptation and visual search in medical images (abstract). J Vision, in press.