Probing the functions of contextual modulation by adapting images rather than observers

Michael A. Webster

Department of Psychology/296, University of Nevada, Reno, Reno, NV 89557, USA

ABSTRACT

Countless visual aftereffects have illustrated how visual sensitivity and perception can be biased by adaptation to the recent temporal context. This contextual modulation has been proposed to serve a variety of functions, but the actual benefits of adaptation remain uncertain. We describe an approach we have recently developed for exploring these benefits by adapting images instead of observers, to simulate how images should appear under theoretically optimal states of adaptation. This allows the long-term consequences of adaptation to be evaluated in ways that are difficult to probe by adapting observers, and provides a common framework for understanding how visual coding changes when the environment or the observer changes, or for evaluating how the effects of temporal context depend on different models of visual coding or the adaptation processes. The approach is illustrated for the specific case of adaptation to color, for which the initial neural coding and adaptation processes are relatively well understood, but can in principle be applied to examine the consequences of adaptation for any stimulus dimension. A simple calibration that adjusts each neuron's sensitivity according to the stimulus level it is exposed to is sufficient to normalize visual coding and generate a host of benefits, from increased efficiency to perceptual constancy to enhanced discrimination. This temporal normalization may also provide an important precursor for the effective operation of contextual mechanisms operating across space or feature dimensions. To the extent that the effects of adaptation can be predicted, images from new environments could be "pre-adapted" to match them to the observer, eliminating the need for observers to adapt.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

While this feature issue is primarily devoted to the functions of spatial context, in the present article I approach the question of contextual modulation by focusing instead on temporal context – on processes that shape a neuron's response that are outside the classical receptive field in time rather than space. The effects of context show a number of parallels between space and time, and appear to modulate sensitivity and perception in similar ways, possibly because the statistics of the visual world are themselves similar in space and time (Schwartz, Hsu, & Dayan, 2007). It would therefore be surprising if spatial and temporal contextual processes were not designed to achieve similar goals. For example, both adaptation and spatial contrast interactions have long been argued to play an important role in color constancy (Brainard & Wandell, 1992; Hurlbert & Wolf, 2004; Land, 1986). A great deal of thinking has been invested in understanding why sensory systems adapt (for recent reviews see (Clifford et al., 2007; Kohn, 2007; Wark, Lundstrom, & Fairhall, 2007; Webster, 2011; Webster & MacLeod, 2011)). Reviewing the functional consequences of temporal context may thus shed light on some of the places to look for analogous roles of spatial context. Alternatively, this can also point to places where they might serve different goals. If temporal context already optimizes some aspects of visual coding, what roles are left for spatial context? And given that both are occurring, how might they interact? To examine these questions I begin by reviewing some of the purposes that have been suggested for adaptation, and then describe a recent approach we have developed for testing different hypothesized functions by "adapting images" to simulate theoretically optimal states of adaptation. This is followed by briefly considering the implications of these temporal processes for mechanisms that respond to spatial context. While the simulations are focused only on the consequences of adaptation for color vision, the principles are general and thus applicable to visual coding in general.

2. The functions of adaptation

Most reviews of visual adaptation include a litany of potential benefits. For example, a typical list might include the following

E-mail address: mwebster@unr.edu

functions of adaptation, along with the ways in which these might be manifest in our perception:

1. *Sensitivity regulation:* Neurons have a very limited dynamic range of their response but must operate over a potentially enormous range of inputs. Adaptation adjusts sensitivity so that the responses are centered around the average stimulus level in the environment (Barlow, 1972). This allows the visual system to devote the full signaling capacity to registering small changes around the mean, where most of the information is concentrated. Behaviorally this shows up as a better ability to distinguish or discriminate stimuli within a scene when we are first adapted to the scene.

2. *Coding efficiency:* The limited capacity of neurons to carry information also exerts strong pressure to optimize efficient coding (Clifford et al., 2007; Wainwright, 1999; Wark, Lundstrom, & Fairhall, 2007). This can act at the level of an individual neuron to adjust its operating curve so that each response level is given equal weight (e.g. histogram matching) (Laughlin, 1981). It can also operate across neurons to remove redundancies in their responses. Adaptation has been proposed to play a role in both of these adjustments. For example, adaptation may adjust the gain of the neural response so that the average response occurs for the average level of the current stimulus gamut (Ohzawa, Sclar, & Freeman, 1982; Rieke & Rudd, 2009); and may potentially remove the correlations between neurons coding different attributes of the stimulus if those attributes are themselves correlated (Barlow, 1990b; Carandini et al., 1997). Signs of these benefits are typically searched for in the neural code rather than the performance of the observer, but should again reflect behavioral benefits in sensitivity and discrimination.

3. *Error correction:* Neurons might also be thought of as mismatched to the stimulus if they signal the "wrong" percept – if the world consistently looks tilted or too yellow. In this case adaptation might serve to recalibrate the neural response to remove the error (Andrews, 1967). For instance, if there are inherent differences in sensitivity to different orientations, then this might introduce biases in the population code for orientation. Adaptation can remove these biases by equating the responses and thus leveling the playing field. Perceptual signs of these corrections are seen in the stability of some percepts despite large sensitivity differences within or between observers.

4. *Perceptual constancy:* Often the visual system is trying to estimate invariant properties of the world from a retinal image in which multiple sources of stimulus variation are confounded. A classic example is color constancy, where the goal is to recover the reflectance of a surface from a spectrum that also varies with the lighting (Foster, 2011; Smithson, 2005). Adaptation can promote constancy by filtering out or discounting some of the sources of variation. Thus when the lighting becomes redder, adapting to become less sensitive to red can remove much of the color shift owing to the illuminant. Here we expect to see advantages of adaptation when we can recognize the same stimulus under different viewing contexts.

5. *Learning and predictive coding:* An efficient way to represent information is to build a prediction about the expected properties of the world and then signal only the deviations from this prediction (Srinivasan, Laughlin, & Dubs, 1982). In this way the generic state of the world need only be represented implicitly. Adaptation could adjust visual coding to null the responses to the expected level. Mechanistically, this should show up as stronger responses the more novel or unexpected the stimulus is (Ranganath & Rainer, 2003). In turn, perceptually this should be manifest so that novel stimuli are more salient, and so that what we notice most about the world are the very properties we are not adapted to.

This list is not exhaustive, but again gives a flavor of the variety of roles that adaptation has been hypothesized to play in sensory systems. Moreover, many of these putative roles are closely related, and may amount to different perspectives on the same problem. However, several issues continue to plague our understanding of the actual purpose of adaptation. First, the wide variety of proposed functions appears to run counter to the fact that, in many cases, it may only be a single type of adjustment that the visual system is making. How can one trick yield so many and so seemingly disparate advantages? This suggests that there is probably some more fundamental and general principle guiding why the visual system adapts, and what we now consider the functions of adaptation will turn out to be specific consequences of this principle. Here I explore the possibility that this principle is "normalization," and reflects the simple tendency of adaptation to balance the responses across the mechanisms coding a stimulus dimension so that they are either nulled or equated within the visual context we are currently adapted to. In the following I use normalization specifically to refer to a form of adaptation, though importantly the term has also taken on a broader meaning which includes both temporal and spatial calibrations (Carandini & Heeger, 2011). Norms are central to models of visual coding (Webster, 2011). In some cases the visual system uses an explicit "norm-based" code in which stimuli are represented relative to a reference or norm which itself is encoded by a null in the neural response and which has a special, "neutral" status in visual coding. Examples of such codes include color vision (where all hues are referenced to gray) and face perception (where all identities are referenced to the average or prototypical face) (Webster & MacLeod, 2011). In other cases, stimuli may be represented by a population code – by something like the peak response in a distribution of neurons tuned for different levels of the stimulus dimensions. Examples in this case include the encoding of orientation or spatial scale. Here the norm is implicit but corresponds to equal activity (no peak) across the set of mechanisms.

The function of adaptation, quite simply, may be to set these norms according to the current context, and can be accomplished simply by adjusting each neuron's gain so that the mean response occurs for the mean stimulus level each neuron is exposed to. There are likely to be many additional forms of adaptation, e.g. to match more complex characteristics of the input (Gollisch & Meister, 2010), or to decorrelate the responses of different neurons (Barlow, 1990b), but here I consider what a visual system can do with only a simple gain change. This modulation alone would assure in theory that the responses of all mechanisms would be the same on average for the current adapting context (i.e. the stimulus distribution we are adapted to). Moreover, this recalibration proceeds naturally whenever the context changes, and thus will always lead to the appropriate norms for the appropriate context. Thus norms are not distinct from adaptation, but rather are synonymous with the states of adaptation that the visual system is currently in. In this sense, the phrase "contextual modulation" is an understatement – context does not merely perturb neural responses, it defines them.

As we will see, it turns out that this normalization predicts each of the functional benefits highlighted above. Yet a second general problem in understanding the purpose of adaptation has been that the behavioral correlates of these functional improvements are often lacking. Strong visual aftereffects can be readily induced for most visual patterns, yet adapting to these patterns often fails to improve visual performance. Specifically, observers are not typically better at detecting or discriminating patterns after they have adapted to them (even though that adaptation leads to large changes in the appearance of the patterns) (Clifford et al., 2007). The clearest exception is light adaptation, where adjusting to the

mean illumination can mean the difference between vision and blindness (Rieke & Rudd, 2009). Here the stimulus range we must operate over is truly daunting, and thus rapid recalibration is essential. Other aspects of the world may vary much less, and thus may not require the same urgency. Yet clearly it is simple to introduce a stimulus that can alter the state of adaptation, so why are the full complement of consequences lacking? Here I examine whether they do in fact become manifest – if observers are given enough time to adapt.

The issue of time reflects a third important area where our understanding of temporal context and its function remains poor. Most studies of visual aftereffects have explored timescales extending at most to a few hours. Yet there is growing evidence and theoretical arguments for sensitivity adjustments over much longer durations (Bao & Engel, 2012; Delahunt et al., 2004; Kording, Tenenbaum, & Shadmehr, 2007; Kwon et al., 2009; Neitz et al., 2002; Shadmehr, Smith, & Krakauer, 2010). Contexts that vary slowly may be optimally tracked by adjustments that themselves vary gradually but lay down more permanent imprints. Indeed, much of the most important recalibrations the visual system faces may reflect changes not in the world but in the observer as they develop and age, and these involve very slow changes in context. This raises the questions of how far adaptation could go in theory to recalibrate norms if given enough time. What would the world look like, and what could we do that we could not before?

3. Simulating adaptation

In this paper I describe an approach we have begun to use for exploring the functional consequences of adaptation, by "adapting images" to simulate how they might appear to an observer under theoretically optimal states of adaptation. To the extent that we understand how a visual attribute is encoded, and how adaptation adjusts this coding, then it is possible to characterize the extent to which percepts could in principle vary across different contexts, or when observers with different visual systems are placed in the same context. The chief advantage of this approach is that it allows one to push the processes of adaptation to their theoretical limit, and thus explore complete states of adaptation that are difficult to probe at short timescales. In turn, these images can be combined with empirical measurements to ask what observers see, and what visual tasks they can perform, when viewing the same contexts under different adapted states.

The theoretical limits of adaptation depend on a number of factors. One is the extent to which the world itself varies. The larger the stimulus variations from one visual context to the next, the more likely it is that adaptation will be important for fine-tuning perception for the specific niche the observer is in. One can explore this by characterizing the range of adapted states predicted by variations in the natural world, and also ask how well we can in principle see within the "unnatural" visual worlds we increasingly experience through technology. A second factor is how much observers vary. Individual differences in visual sensitivity begin as early as the optics but arise continuously throughout the visual stream (Wilmer, 2008). One can ask how the adaptation adjusts for these variations and thus how important they ultimately are for different perceptual tasks. In fact an advantage of the present approach is that it provides a common framework for evaluating the influence of environmental vs. observer variations on perception. A third factor is the actual structure of the visual representation (e.g. the number or nature of the presumed visual channels coding a given stimulus dimension). This will necessarily limit the number and types of ways that the visual system can adjust to the stimulus. Examining these adjustments for a presumed model of visual coding can reveal how well that particular neural

architecture could be optimized for different contexts, and thus which specific attributes of the stimulus it might be designed to be adaptable to. Moreover, if it turns out the model performs less well than real observers, then that can point to either deficiencies in the presumed neural model or to additional forms of plasticity. Finally, the limits will also depend on the nature of the adaptation itself. For example, adaptation could alter the responses of neurons in a variety of ways, and also potentially alter the interactions between neurons. The specific form of these adjustments may lead to very different end states, and exploring these and the functional advantages they offer can shed light on the processes of adaptation.

4. Adaptation and color coding

To examine these questions I have restricted the analysis to the problem of color adaptation. Compared to other stimulus domains, the early stages of color coding and how these are affected by adaptation are well understood (Webster, 1996). Moreover, the form and basis of individual differences in spectral sensitivity are well known, and a number of studies have documented how color varies across different natural environments (Lee, Wachtler, & Sejnowski, 2002; Long, Yang, & Purves, 2006; Ruderman, Cronin, & Chiao, 1998; Tkacik et al., 2011; Webster, Mizokami, & Webster, 2007; Webster & Mollon, 1997). Thus color is a case where reasonable predictions can be made about the consequences of adaptation. As noted however, the principles themselves are nevertheless general, and hopefully illustrate the general roles of adaptation on perception.

Predicting how the visual system adapts to color again requires information about the stimulus and the observer. This includes the nature of the stimulus context (e.g. the distribution of colors the observer is exposed to); the number and sensitivity of the mechanisms that encode color; and the ways in which adaptation adjusts their responses. Recently we developed a model that incorporates these factors to simulate how the colors within images should appear when observers are under different states of chromatic adaptation (Juricevic & Webster, 2009; McDermott et al., 2008; Webster & Juricevic, 2013; Webster, Juricevic, & McDermott, 2010). The model is simple but based on plausible assumptions about color coding and adaptation (Fig. 1). At the initial level, light at each point in the image is absorbed by the L, M, and S cones. The signals from these cones are then combined to form an array of postreceptoral mechanisms that sample different directions within the sphere of color space. The three channels shown in the figure correspond to the three cardinal mechanisms along which chromatic and luminance information are thought to carried in the retina and LGN (Derrington, Krauskopf, & Lennie, 1984). However, for the model we instead included 26 mechanisms that spanned the cardinal axis space in 45 deg intervals. These were formed by different linear combinations of the cardinal axis mechanisms, so that each channel's sensitivity varied as the cosine of its preferred direction in the volume of color space. The large number of channels was included because chromatic contrast adaptation can itself be selective for any arbitrary direction in color space, implying that the space is effectively tiled by a dense array of mechanisms each tuned to a different chromatic direction (Krauskopf et al., 1986; Webster & Mollon, 1994). (Variations in the specific number do not markedly affect the predictions, though reducing the set to three again fails to capture the observed selectivity of the adaptation.) Finally, color at each point is based on summing the outputs across the mechanisms projected onto the cardinal axes, with the responses initially normalized so that this returned the original image colors for the reference state.

Adaptation within the model occurs at both stages. In the receptors, the cone sensitivities are multiplicatively scaled so that the
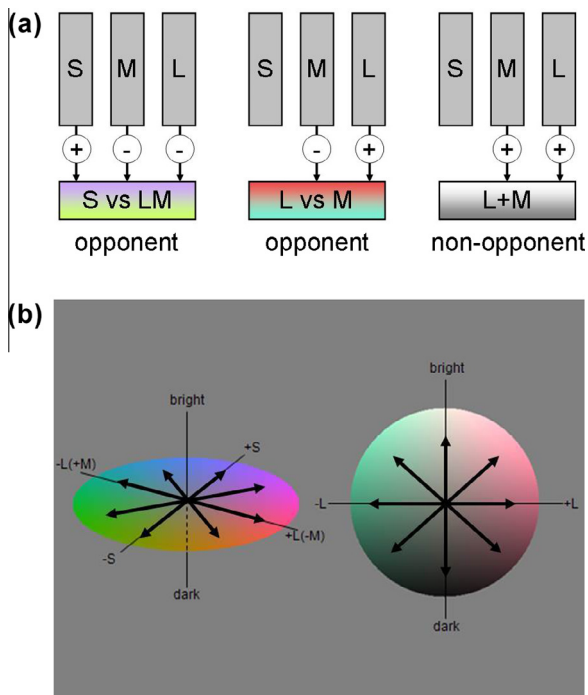
Fig. 1. Model of color coding and adaptation. Adaptation adjusts the responses at two stages (corresponding to the cones and postreceptoral channels) so that within each mechanism the average response to a new context equals the average response to the reference context. Postreceptoral channels included 26 mechanisms formed by different linear combinations of the cones. These include the three cardinal mechanisms (SvsLM, LvsM, or L + M), and additional "higher-order" channels tuned to axes at intervals of 45 deg in the cardinal axis space.

average response to the current environment within each cone equals the average response to the reference environment. This adjustment is known as von Kries scaling and is a well-established characteristic of chromatic adaptation (Brainard & Wandell, 1992; Chichilnisky & Wandell, 1995; Webster & Wilson, 2000; Wuerger, 1996). In the postreceptoral mechanisms, the contrast gains are again scaled so that within each the average response to chromatic contrast in the current environment equals the channel's response in the reference environment. This multiplicative normalization is not in fact characteristic of short-term contrast adaptation measured psychophysically, which instead more closely follows a subtractive change (Georgeson, 1985; Webster & Mollon, 1994). However, there is evidence for shifts to a more multiplicative or response gain change at longer durations (Kwon et al., 2009), and this form of adaptation is necessary to normalize the channels. To simulate adaptation, the responses within each mechanism are averaged for some ensemble color distribution representing a reference or baseline state, and then for a distribution corresponding to the current context. The responses to the latter are then rescaled so that the average response within the current context is the same as the reference context, and the image colors are then re-rendered from the adapted outputs. As discussed below, the reference context itself can be a color distribution from a specific image or environment, or a theoretical distribution (e.g. predicted from a uniform color space). For the simulations shown here we used the former, so that the adaptation effects predict how different the colors under adaptation to a changed environment or spectral sensitivity should appear relative to your percepts of the original image.

Note that even in this very simple model adaptation is adjusting the norms at a number of stages. In the cones, the adaptation levels the responses for the average chromaticity in the scene and thus represents gray by equal activity across the receptors. At the

postreceptoral level, this gray level is instead set to correspond to the null point across the opponent channels. Finally, the model assumes that hue is represented by a population code, and adaptation equates the average activity levels across the population.

5. Color constancy and compensation for variations in the observer's sensitivity

As noted, one of the putative roles of adaptation is to promote stable percepts by filtering out extraneous variations in the stimulus. In color this has been intensely investigated when the extraneous change is the lighting (Foster, 2011; Shevell & Kingdom, 2008), but also routinely arises because of sensitivity variations in the observer (Werner, 1996). The latter has also been examined in many studies, and provides a powerful natural experiment for exploring the consequences of very long-term adaptation and the extent to which this compensates for the idiosyncracies of an individual's eyes. For example, the density of lens pigment steadily increases with age, and thus less and less short wavelength light reaches the retina. However, the color appearance of the world remains highly stable with aging (Werner & Schefrin, 1993). Similarly, observers vary dramatically in the relative numbers of the different classes of cones, yet these differences fail to predict differences in color appearance (Brainard et al., 2000; Miyahara et al., 1998). In both cases this stability has been attributed to long-term adaptation, and falls readily out of the simulations of adaptation. Recently we modeled the predicted changes in color appearance from variations in optical screening pigments or color deficiencies (Webster et al., 2010). Von Kries adaptation (receptor gain changes) is sufficient to discount almost all of the color shift produced by differences in lens or macular pigment screening (Fig. 2). Similarly, when the model is used to simulate the percepts of anomalous trichromats, their color experience is much more similar to color normals than their sensitivity would suggest (Fig. 3). In this case the reason is because in addition to receptor scaling, postreceptoral channels are adjusting their gain according to the gamut of the receptor inputs. An anomalous trichromat has two cones that have closely similar spectral peaks, and thus the difference signal carried by these cones is necessarily reduced. Yet if the cone contrasts are reduced then adaptation is predicted to amplify the gain in post-receptoral channels, again largely discounting this effect. Importantly, something like this actually seems to occur in some color deficients – for some the world does not appear as reduced in color as their cone sensitivities might suggest (MacLeod, 2003; Regan & Mollon, 1997). It is also something that happens in all our eyes. The cone differences signaling color are many times smaller than the sums on which luminance depends. Yet the world does not appear impoverished in color, presumably because the gains of postreceptoral mechanisms are each adjusted for their respective gamut. In fact the world looks about as colorful as one would predict from adaptation to the range of color signals in the world (McDermott & Webster, 2012a).

There are a number of important points to note with regard to how adaptation promotes color constancy by adjusting for the sensitivity limits of the observer. First, these adjustments do reflect long-term calibrations. This can be seen in cataract patients who take weeks or months to recalibrate their sensitivity after their lens is replaced (Delahunt et al., 2004). This differs from the rapid von Kries adaptation that is normally studied in color constancy, which can itself quickly recalibrate the color norm. Two observers with different densities of lens pigment will describe the same stimulus as white if they look at the same stimulus for a few moments – because the visual system will rapidly adapt to the current input. However, the stability of white also holds for observers while dark adapted and thus not under the spell a specific stimulus. The site of both

Fig. 2. Predicted consequences of an aging lens. Images show the world as it might appear to an observer with the average lens pigment density of a 13 year old (left), the same image filtered through the eyes of an observer with the average density of a 70 year old (center), or the image as it might appear if the visual system is adapted to the lens density change. The predicted color shifts are also illustrated for an image of the Munsell palette of colors (from Webster, Juricevic, and McDermott (2010)).
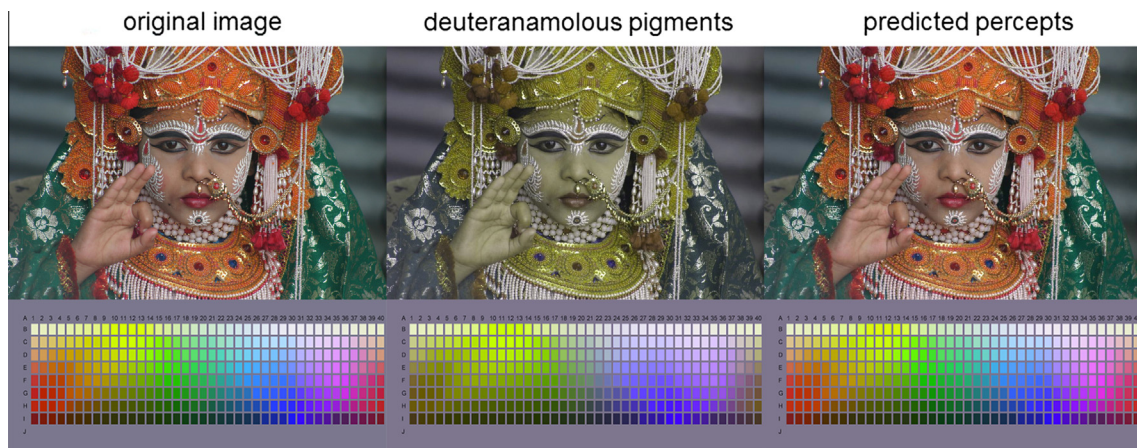

Fig. 3. Compensating color appearance for a color deficiency. Chromatic signals in a deuteranomolous observer are strongly reduced by the close spectral peaks of their longer-wave cones (center image). Adaptation may restore the perceptual gamut of color, even if this also amplifies the noise and thus does not enhance discrimination (right image). Palettes below each image show the predicted color changes within the Munsell palette of colors (from Webster, Juricevic, and McDermott (2010)).

forms of adaptation can be localized to the receptors themselves, and thus reveals that the cone sensitivities are set by mechanisms operating over multiple timescales (Webster & Leonard, 2008).

A second point is that the same processes not only adjust percepts for the sensitivity differences between observers, but more importantly adjust for differences within the observer. There are enormous variations in chromatic sensitivity with eccentricity, and these do lead to varying hue percepts in the periphery (Abramov & Gordon, 1994). However, the visual field appears fairly constant in its colorfulness – and does not obviously shift in color whether we are looking through or outside the regions filtered by macular pigment – and at least part of this can be traced to spatially local processes of adaptation (O'Neil & Webster, 2014; Webster & Leonard, 2008).

Third, color percepts are in fact more stable than predicted by multiplicative gain changes in the receptors, which can discount

most but not all of the spectral shifts from an illuminant change or a change in prereceptoral screening pigments (Bompas, Powell, & Sumner, 2013; O'Neil & Webster, 2014; Webster et al., 2010). This illustrates a case where modeling specific forms of adaptation can be used to reveal the presence of additional compensatory processes. (The studies of Bompas, Powell, and Sumner (2013) and O'Neil and Webster (2014) also draw attention to the finding that this compensation is not perfect, yet the residual errors reflect very minor deviations relative to the enormous degree of constancy already afforded by von Kries scaling, which again underestimates the level of constancy observed.)

Fourth, many attempts have been made to simulate how the world appears to someone with a different visual system – to a color deficient adult or a color normal baby. Yet most of these are based only on filtering the stimulus according to the subject's spectral sensitivity. This accurately portrays the information avail-

able to them, but probably misrepresents how that information is experienced, for these individuals will also adapt to compensate for their sensitivity.

The final point is what specifically the adaptation predicts about the way we experience color. The fact that lens pigment density has little impact on how observers describe colors shows that color percepts are not strongly constrained by our spectral sensitivity, but in turn the adaptation itself also places little constraint on those percepts. Specifically, there are enormous individual differences in how people judge color, such as the spectra corresponding to unique hues (pure red, green, blue or yellow) (Kuehni, 2004). It is unlikely that these result because observers are in different adaptation states. Instead, these differences reflect the general problem that we do not yet know the neural basis for color percepts, and thus cannot in general predict which stimulus an observer might select for a given color. However, an important exception is the stimulus that appears white. Here this norm is arguably set by the observer's adaptation, and whether two observers experience the same stimulus as white can plausibly be related to the history of colors they have been exposed to. Thus norms may represent a special case where we can relate perceptual experience to a special state in the underlying neural code (Webster & Leonard, 2008). In calibrating these norms for a common visual environment adaptation could serve an additional function in contributing to a form of shared consciousness or "interobserver constancy" that might allow us enough phenomenal similarities in our visual experience that we can meaningfully communicate with each other about it (Webster, Werner, & Field, 2005a).

6. Adaptation to specific environments

To summarize, despite marked differences in sensitivity both between and within observers, to the extent that they are adapted to a common environment their percepts should tend to converge, and in this regard adaptation plays a fundamental role in compensating for these sensitivity differences to promote color constancy. But what about two observers adapted to different environments? Does the world itself vary enough in its color characteristics that color perception might vary depending on the specific natural context we are in? To examine this, we can simulate how large the changes in color percepts might be within the same observer if they were completely adapted to one environment vs. another. Fig. 4 provides an illustration of these simulations (Juricevic & Webster, 2009). The upper panels show images from a valley in the Western Ghats of India during the monsoon or winter, and illustrate the often large seasonal variations in the color distributions within the same environment. The lower panels show the same images after adapting each image such that the average response of each model neuron when exposed to the current environment is the same as when exposed to the alternative environment. Specifically, the lower left (right) image shows how the monsoon (winter) scene might appear to an observer adapted to the monsoon (winter) season. The most obvious effects are that the dominant physical chromaticities within the scene (e.g. the lush greens in the rainy season) become muted or desaturated, while more novel colors in the scene become highlighted. For example, during the arid season green foliage is rare, and these greens are predicted to become more vivid through adaptation to the arid color distribution.

One implication of this illustration is that even natural and routine variations in the color of the world are potentially large enough to place people within different worlds of color experience. The apparent color of the same surface, and which surfaces stand out, might be significantly different for two individuals, and this is not simply a matter of learned familiarity or novelty but is a form

of learning and novelty that derives explicitly from how adaptation adjusts the visual representation of color. That is, in a dry world, the greens really do _look_ greener. This also suggests that were we to sample the color percepts of people living in different places in the world, then we might expect their percepts to vary in part because the color of the world itself varies.

Population differences in color perception have in fact been widely investigated and are at the center of the on-going debate over whether color percepts are universal or shaped by culture and language (Kay & Regier, 2006). But to what extent could these differences also be shaped by the environment? Figs. 5 and 6 show a previously unpublished analysis of this question. Imagine using the Munsell palette in Fig. 5 to choose the chips that represented the best examples of different color terms such as red, green, or blue. We can then model how these choices might shift when the same selections are made while adapted to different environments in the world. To examine this, color distributions were collected from a range of different settings. These included calibrated datasets of natural scenes in the Western Ghats and Sierra Nevada (Webster, Mizokami, & Webster, 2007), which are likely to be typical of many natural outdoor environments. They also included uncalibrated image sets by sampling the internet for scenes of "characteristic" natural environments (e.g. forests or meadows), or environments that were "uncharacteristic" (e.g. tundra or desert), or "unnatural" (outdoor suburban settings). For each of these environments, the model visual system was adapted to the appropriate color distribution, and then the palette chip required to give the same response criterion for a given focal color was calculated from the adapted channels.

These values are shown in Fig. 6, which plots how much an observer's choice of red, green, blue or yellow would be predicted to vary if we shifted the observer between the different environments. For comparison, the lighter symbols plot actual measurements of population differences in focal color choices from the World Color Survey (Webster & Kay, 2007), which sampled color naming in respondents from 110 languages scattered throughout the world (Kay et al., 2009). The mean focal choices for different linguistic groups significantly vary, but nevertheless show pronounced clustering around common regions of the space. The pattern for the simulated choices show a number of striking parallels. First, there are clear differences in the choices owing to differences in the state of adaptation. Second, the spread is roughly of the same order of magnitude as the empirically measured choices. Third, for both datasets, the effects of the population or modeled differences is restricted to biasing the location of the focal stimuli to different locations within the color category. That is, different linguistic groups, and different visual worlds, do not differ enough so that what signals "red" in one case can shift to "yellow" in another (at least as far as these constructs can be measured). Of course, this is at best a thought exercise, since we do not yet know enough about how the color statistics of the world actually vary, let alone for the WCS respondents. Nevertheless, it suggests that differences in visual contexts alone are sufficient in theory to account for the range of differences in average color choices that have been observed across different cultures.

7. Adaptation, visual salience, and attention

If we return to Fig. 4, note again that the primary effect of adaptation is to tone down the dominant colors while drawing out the more novel colors. This closely parallels predictive coding since the expected stimulus features are effectively nulled. This coding is also more efficient and metabolically-effective (Lennie, 2003) because it represents the dominant properties of scenes only implicitly. The ways in which adaptation recodes the stimulus in
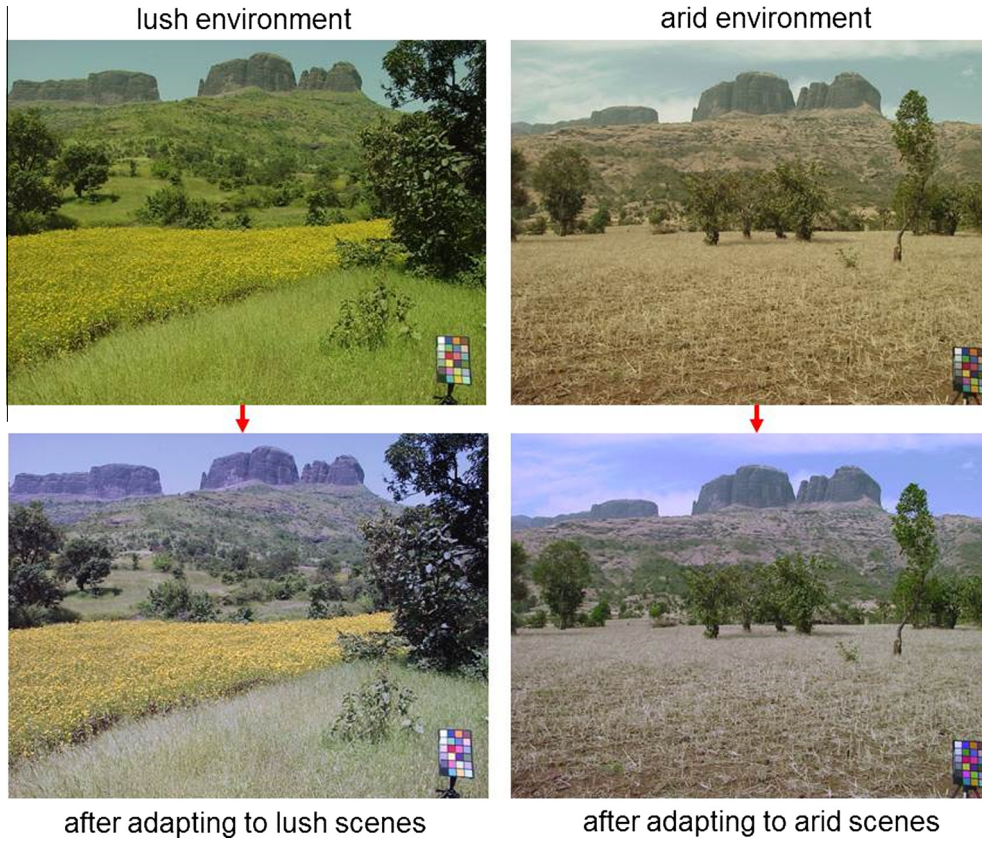
Fig. 4. Perceptual color shifts predicted by adaptation to a lush or arid environment. The top images show scenes from each environment as viewed under a common state of adaptation. The lower images show the perceived colors predicted for individuals adapted to each environment (from Webster (2011)).
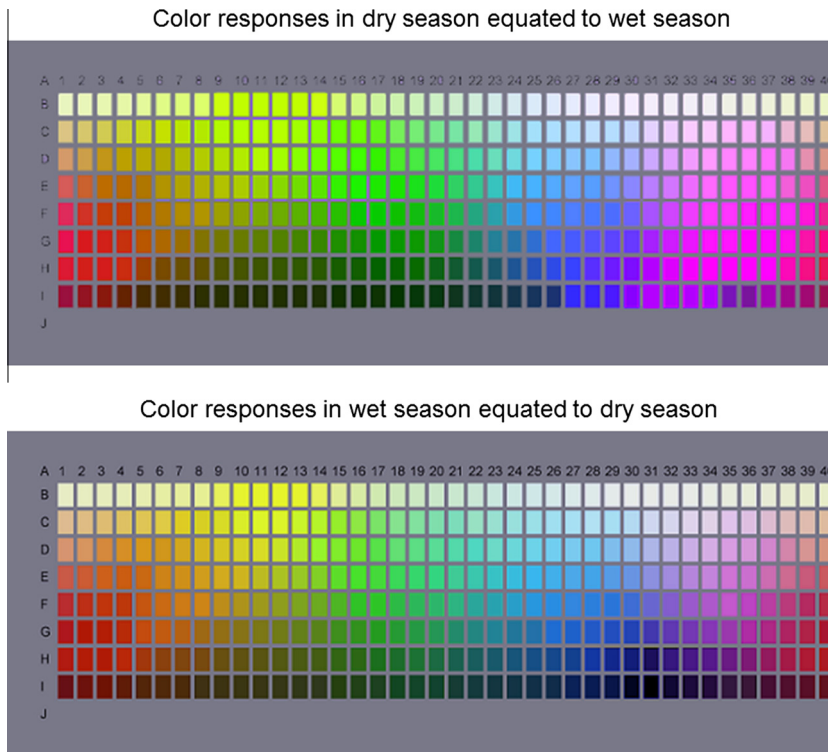

Fig. 5. Predicted shifts in the hues of the Munsell palette between observers adapted to a lush or arid environment. For each, the palettes simulate how the hue of each chip should appear to an observer adapted to one environment (e.g. dry season) so that the average response within each mechanism equals the average response under the alternative environment (e.g. wet season).
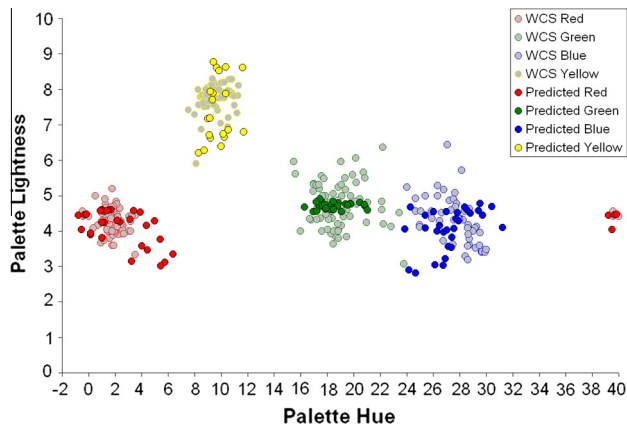
Fig. 6. Range of red, yellow, green, or blue hue settings predicted for adaptation to different environments (color symbols) compared to observed range of variation across different populations in the World Color Survey (gray symbols). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

terms of deviations from the norm may impact what we notice about the world. To an explorer arriving on Mars all rocks may start out as gems, but over time it will be the rarer elements that start to catch the eye. Recent studies have revealed a close connection between normalization processes and the actual mechanisms of attention (Reynolds & Heeger, 2009). Specifically, attention increases the gain of the neurons prior to the normalization. Adaptation also provides a bottom-up influence on attention, by modulating the relative salience of stimuli such that the unexpected features become more conspicuous (Barlow, 1990a, 1990b).

These increases in salience predict that novel stimuli or "statistical outliers" should be more readily detectable after adapting to an environment. However, as noted, one of the hurdles to understanding why vision adapts has been in demonstrating that visual performance actually improves following adaptation to different kinds of patterns. Specifically, beyond retinal light (e.g. Rieke & Rudd, 2009) and chromatic (e.g. Krauskopf & Gegenfurtner, 1992) adaptation, the large changes that adaptation induces in appearance are often not accompanied by measurable changes in the speed or accuracy of visual judgments. This could be because normalization does not impact sensitivity for most stimuli, or perhaps because the world does not vary enough in these stimuli so that sensitivity adjustments are necessary. Alternatively, it could again be the case that these sensitivity adjustments are fundamentally important, but do not become strongly manifest over the short durations typically used to measure adaptation effects. Again, to explore this we can ask how well we could detect information within an environment if we were under theoretically complete adaptation to that environment (Webster & Juricevic, 2013).

Fig. 7 examines this first for two extreme environments – under water or on the surface of Mars. The images were adapted so that within each the average mechanism outputs equaled the average responses to an arid terrestrial environment on Earth. Note again that these are the same images except for the presumed state of the adaptation. Clearly, adaptation to each specific environment reveals a great deal of structure in the images that is not available to the incorrectly adapted eye. In particular, the relatively low chromatic contrasts of both environments are amplified by the adaptation.

To empirically test changes in visual sensitivity, a visual search task was used where the observer had to identify the location of a target color within a background of distractors, both superimposed as 10 a by 10 array of elements added onto the background images,
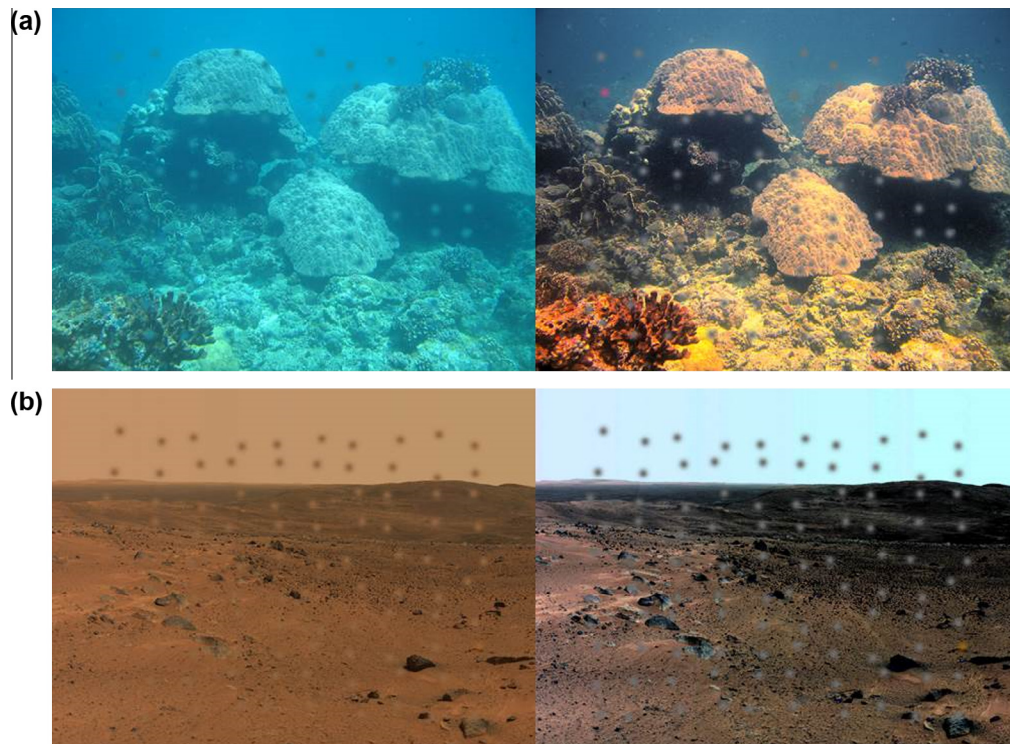

Fig. 7. Color shifts predicted from adaptation to an underwater environment or the surface of Mars. Superimposed blobs show arrays of a target and gray distractors used to measure search times for detecting the target in the original or adapted images (from Webster and Juricevic (2013)). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

as shown in Fig. 7 (Webster & Juricevic, 2013). The distractors all had the mean chromaticity of the background, and thus were gray within the adapted images. The color differences between the targets and distractors were quickly evident in the adapted images, while often imperceptible in the original images, and thus visual search was far superior for the adapted images (Fig. 8). (In fact this figure underestimates the size of the effects, since many target colors in the original images could not be located within the 5 s time limit allowed for each search.) These results for simulated adaptation states are consistent with recent empirical studies that have also revealed improvements in visual search when observers are first adapted to the backgrounds they are searching within (Kompaniez et al., 2013; McDermott et al., 2010; Wissig, Patterson, & Kohn, 2013).

The extreme cases in Fig. 7 might seem a poor test of the function of adaptation in optimizing color salience and discrimination, except when considering that how we are calibrated for very different worlds may be no worse than how we were initially calibrated for our own. That is, it seems plausible to assume that the close match of vision to our own environment required initially very large adaptational adjustments of the kind illustrated in Fig. 7, and that it is only because of these that the gamut of colors in our world seems as rich and balanced as it does.

It remains less certain whether adaptation is necessary to optimize sensitivity for the specific natural environments that we occupy. However, one hint of this is illustrated in Fig. 9 (Webster & Juricevic, 2013). In this case observers searched for targets within lush environments shown under simulated adaptation to that environment or to an arid color distribution more typical of the vicinity of Reno, NV where the observers lived. This again led to significant improvements in search times for some target colors, suggesting that this routine variation in the natural world is sufficiently large so that visual performance can be optimized for specific natural contexts.

8. Adaptation and visual discriminations

As illustrated above, one predicted effect of adaptation is to expand the perceptual gamut of contrasts when these are low, either in the stimulus (e.g. Fig. 7) or because of properties of the observer (Fig. 3). These contrast gains accentuate the differences between stimuli and thus should enhance not only suprathreshold detection, as measured in the search task, but also threshold discriminations (though whether increases in contrast sensitivity benefit discrimination depends in part on the site of the limiting noise, as noted below (Rieke & Rudd, 2009; von der Twer & MacLeod, 2001)). In the case of brightness or color, centering the mean response at zero contrast (gray) further enhances threshold discriminations around the adaptation point (Craik, 1938; De Valois & Switkes, 1983; Krauskopf & Gegenfurtner, 1992).

9. Limitations of the model

The preceding examples illustrate that a very simple calibration of neural responses affords a host of benefits, for both perception and performance and for optimizing efficient coding. Are these simulations plausible? There are number of ways in which the model we used is highly simplistic. For example, we sampled color values evenly and independently, within linear chromatic mechanisms with no spatial extent, and assumed multiplicative gain changes when these have not been well established for contrast adaptation. A further omission, which may seem scandalous given the theme of this issue, is that we treated the responses of the channels as independent and thus did not include cross-channel normalization, or the effects of spatial context.

More critically, the model we used did not incorporate noise. If adaptation acts after the limiting source of noise then it might have large influences on appearance but little benefit for sensitivity (Rieke & Rudd, 2009; von der Twer & MacLeod, 2001). Again, an argument against this early noise is the relative sensitivity to luminance and chromatic contrast, which is effectively compensated for the large differences in cone contrasts, though the evidence
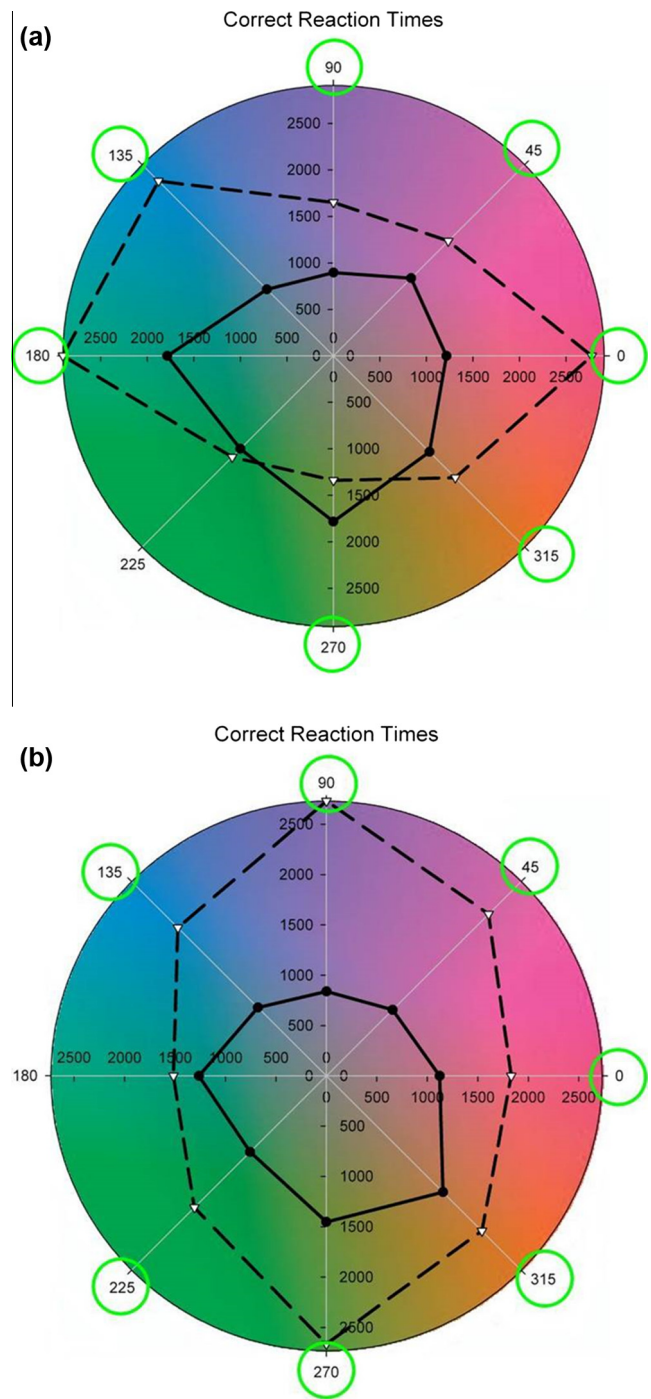

Fig. 8. Mean reaction times for detecting the target colors in the original (dashed lines) or adapted images (solid lines). Targets were defined by hue angles at 45 deg intervals in the LM vs. S plane. Circled angles correspond to differences that were statistically significant. Left: underwater images; right: Mars images (from Webster and Juricevic (2013)). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Fig. 9. Search times within lush scenes before (dashed) or after (solid) adaptation to a lush environment. Targets were defined by hue angles at 45 deg intervals in the LM vs. S plane. Circled angles correspond to differences that were statistically significant (from Webster and Juricevic (2013)). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
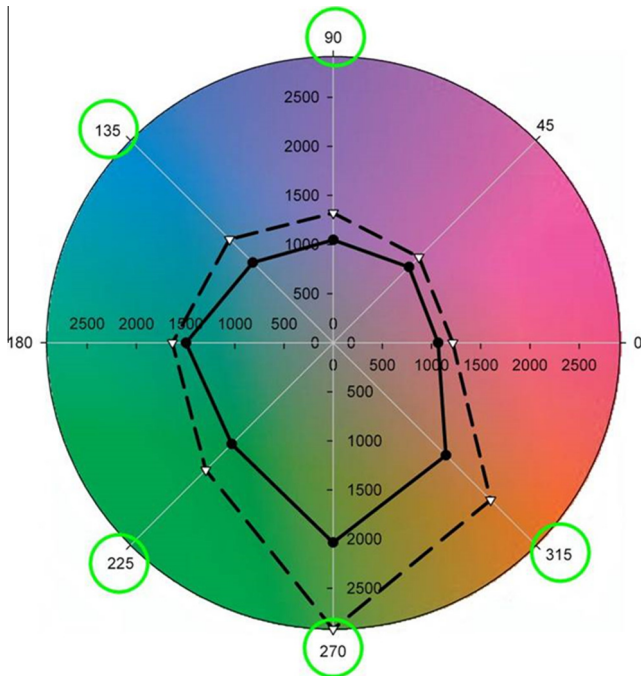
for scaling signal vs. noise is more ambiguous for color deficients (Regan & Mollon, 1997).

A further issue for the simulations is deciding what represents an appropriate reference environment. In principle, this reference is not needed, since the point of the adaptation is to sphere the mechanisms responses for the current environment. However, this depends on assumptions about the mechanisms, and the reference was included to ensure that the model returned the original colors for the original adaptation state. For predicting large changes in the environment, reasonable approximations can be made about the environment we are normally adapted to. However, these approximations become less certain as the environments themselves become more similar. Interestingly, for color vision there are theoretical estimates of the color world a standard observer is adapted to. These are provided by uniform color spaces, which try to scale color differences so that equal distances anywhere within the space correspond to equal perceptual differences. This is in principle the perceptual space that adaptation should impose on our color vision, and it is possible to work back from these uniform spaces to ask what stimulus distributions they correspond to (McDermott & Webster, 2012b).

A remaining limit is that we currently know little about the functional implications of the time course of adaptation (Webster, 2011). If these optimizations are so useful, why does the visual system not rapidly renormalize whenever the context changes? Part of the answer is that this would actually impede perception. If we rapidly adapted to each scene then all colors and the scene itself would look gray. In fact some components of adaptation are so rapid that this fading is prevented only because of eye movements (Ditchburn & Ginsborg, 1952). Yet the time constants controlling other components appear surprisingly long (e.g. Delahunt et al., 2004). The number and optimal timescales for contextual modulation and how these may be tied to different perceptual tasks remain an important and unexplored area of visual coding.

10. Limitations of the adaptation

Pushing adaptation to its limits is also instructive for probing what adaptation cannot do. As noted, this will depend on both the structure of the visual system and the form of the response changes induced by adaptation. As an example of the former, the number and selectivity of the channels representing a stimulus dimension may fundamentally limit the ways in which the visual system can adjust to that dimension. Thus three classes of cones severely restricts the ways that color vision can adapt to a spectral change. Similarly, if adaptation is restricted to simple gain adjustments, then this alone would not allow neurons to adjust their responses according to higher-order moments of the stimulus distribution. The concept of metamers has been widely exploited in color vision and to a lesser extent other stimulus domains (Balas, Nakano, & Rosenholtz, 2009; Freeman & Simoncelli, 2011; Richards, 1979; Williams, Tweten, & Sekuler, 1991), to probe the limits of visual coding. In the same way, there exist "adaptation metamers" which represent stimuli which are physically and sometimes also perceptually different but which produce identical adaptation states (Webster, Werner, & Field, 2005b). Characterizing these for different contexts could help reveal both the actual nature of adaptation and which properties of the world the visual system is ultimately designed to be calibrated for.

With regard to its functional consequences, it is also important to emphasize that a single adjustment like normalization cannot be simultaneously optimal for all perceptual goals. For example, Webster and Mollon showed that adaptations to contrast that might improve coding efficiency are not necessarily the ones that will promote color constancy (Webster & Mollon, 1995). Moreover, Abrams, Hillis and Brainard found that discrimination and constancy cannot always be optimized by a single common mechanism (Abrams, Hillis, & Brainard, 2007). A general example of this problem in color vision is the "gray world assumption," where cone-specific adaptation might help factor out the color of the illuminant by adjusting to the average reflectance of a scene (Buchsbaum, 1980). This adjustment discounts the lighting when the average reflectance is flat, but will lead to failures of constancy when the scene itself is biased. In this way, the graying of Mars with time – though likely facilitating color discriminations – can also be viewed as a failure of color constancy. Finally, such failures in part result from the "coding catastrophe," or the fact that the visual system may not know the state of adaptation that it is in, and thus interprets changes in these states as changes in the stimulus (Schwartz, Hsu, & Dayan, 2007; Series, Stocker, & Simoncelli, 2009). This misattribution is in fact what is thought to underlie the seemingly illusory nature of visual aftereffects.

11. From temporal to spatial context

As noted at the outset, there are numerous parallels between contextual effects that occur over time and space, and similar functional accounts have been applied to both (Schwartz, Hsu, & Dayan, 2007), Like adaptation, spatial induction effects arise at a number of stages in the visual system, even for a common attribute like color (Rinner & Gegenfurtner, 2000; Shevell, Holliday, & Whittle, 1992) and can adjust percepts to distinct properties of the stimulus in ways similar to adaptation. Thus both adaptation and induction adjust color appearance according to both the mean and variance of the surround (Brown & MacLeod, 1997; Singer & D'Zmura, 1994; Webster et al., 2002). Both can also adjust to high-level perceptual attributes (Shevell, 2012; Shevell & Kingdom, 2008; Webster, 2011). On the other hand, in the case of lightness and color, spatial context adjustments can include a wealth of phenomena associated with the perceived layout and lighting of scenes and

the material properties of objects (Kingdom, 2011). Whether there are comparable factors affect adaptation to color remain relatively unexplored (e.g. Goddard, Solomon, & Clifford, 2010). Moreover, the form of the color changes produced even by simultaneous contrast in simple patterns remains uncertain and task-dependent (Bosten & Mollon, 2012), so that it is not clear to what extent the adjustments are analogous to the forms of scaling induced by adaptation.

Many spatial contextual effects on neural responses and visual sensitivity can be accurately described by a divisive normalization that pools the signals across a broad range of neurons to set the gain on an individual neurons' response (Carandini & Heeger, 2011). The benefits of these interactions are again thought to parallel many of the functional consequences of temporal adaptation (Carandini & Heeger, 2011; Schwartz, Hsu, & Dayan, 2007). One important difference however is that the spatial or cross-channel normalization may help to sharpen the neural output by producing a winner-take-all response across the population (Busse, Wade, & Carandini, 2009). This is the opposite of adaptation, which instead acts to "punish" the winners if their responses are consistently too strong, a rebalancing designed to equate the responses across channels. A compelling example of this balancing was recently reported by Benucci, Saleem and Carandini, who found that the population responses of V1 cells quickly adjusted to remove biases in response to stimuli with a biased sample of orientations (Benucci, Saleem, & Carandini, 2013). Clearly, equating neural responses seems an important precursor to cross-channel normalization, for if the prior responses are not appropriately balanced then the inherent bias would disproportionately favor the more responsive mechanisms. Thus a final function that can be noted for temporal context is to appropriately calibrate the mechanisms of spatial contextual modulation. This also suggests that the adaptation should set the intrinsic gain of the mechanisms prior to the cross-channel modulation. Simulations of adaptation that incorporate simultaneous contextual interactions could help visualize the interplay of these different processes in shaping our percepts and further reveal their potential functions.

12. Conclusions

To summarize, we have argued that context goes far beyond modulating neural responses and visual coding to fundamentally establishing how information is represented, by defining the norms on which visual coding depends. Simple processes of normalization acting within individual neurons or channels are sufficient to effectively calibrate visual sensitivity. These calibrations generate a wide array of benefits ranging from how things look to how well we can see them. Demonstrating these advantages can sometimes be difficult because we are already adapted and thus optimized for the current environment, and the full benefits of a recalibration may take much longer to unfold than can practically be measured in the laboratory. We describe a way of circumventing this problem by instead adapting images to simulate the consequences of a change in adaptation states. This provides a novel technique for exploring the functional consequences of adaptation by pushing the effects of adaptation to their theoretical limit. It also allows a common framework for understanding and visualizing how perception changes when the environment or the observer changes. Finally, a practical advantage is that it obviates the need for observers to adapt. To the extent that we understand how the color or spatial statistics of the world are encoded by the visual system and how it adapts to changes in these statistics, images from a new environment could be pre-adapted so that they are optimized for the observer.

References

Abramov, I., & Gordon, J. (1994). Color appearance: On seeing red–or yellow, or green, or blue. *Annual Review of Psychology, 45*, 451–485.

Abrams, A. B., Hillis, J. M., & Brainard, D. H. (2007). The relation between color discrimination and color constancy: When is optimal adaptation task dependent? *Neural Computation, 19*(10), 2610–2637.

Andrews, D. P. (1967). Perception of contour orientation in the central fovea Part 1: Short lines. *Vision Research, 7*, 975–997.

Balas, B., Nakano, L., & Rosenholtz, R. (2009). A summary-statistic representation in peripheral vision explains visual crowding. *Journal of Vision, 9*(12), 11–18 (article 13).

Bao, M., & Engel, S. A. (2012). Distinct mechanism for long-term contrast adaptation. *Proceedings of the National Academy of Sciences of the United States of America, 109*(15), 5898–5903.

Barlow, H. B. (1990a). Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. *Vision Research, 30*(11), 1561–1571.

Barlow, H. B. (1990b). A theory about the functional role and synaptic mechanism of visual aftereffects. In C. Blakemore (Ed.), *Visual coding and efficiency* (pp. 363–375). Cambridge: Cambridge University Press.

Barlow, H. B. (1972). Dark and light adaptation: Psychophysics. In D. Jameson & L. M. Hurvich (Eds.). *Handbook of sensory physiology* (Vol. VII/4, pp. 1–28). New York: Springer.

Benucci, A., Saleem, A. B., & Carandini, M. (2013). Adaptation maintains population homeostasis in primary visual cortex. *Nature Neuroscience, 16*(6), 724–729.

Bompas, A., Powell, G., & Sumner, P. (2013). Systematic biases in adult color perception persist despite lifelong information sufficient to calibrate them. *Journal of Vision, 13*(1).

Bosten, J. M., & Mollon, J. D. (2012). Kirschmann's fourth law. *Vision Research, 53*(1), 40–46.

Brainard, D. H., Roorda, A., Yamauchi, Y., Calderone, J. B., Metha, A., Neitz, M., et al. (2000). Functional consequences of the relative numbers of L and M cones. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 17*(3), 607–614.

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: Optics, Image Science, and Vision, 9*(9), 1433–1448.

Brown, R. O., & MacLeod, D. I. (1997). Color appearance depends on the variance of surround colors. *Current Biology, 7*(11), 844–849.

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

Busse, L., Wade, A. R., & Carandini, M. (2009). Representation of concurrent stimuli by population activity in visual cortex. *Neuron, 64*(6), 931–942.

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 Transaction of the Royal Society of London B: Biological Sciences, 352*(1358), 1149–1154.

Carandini, M., & Heeger, D. J. (2011). Normalization as a canonical neural computation. *Nature Reviews Neuroscience, 13*(1), 51–62.

Chichilnisky, E. J., & Wandell, B. A. (1995). Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. *Vision Research, 35*(2), 239–254.

Clifford, C. W., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., et al. (2007). Visual adaptation: Neural, psychological and computational aspects. *Vision Research, 47*(25), 3125–3131.

Craik, K. J. (1938). The effect of adaptation on differential brightness discrimination. *Journal of Physiology, 92*(4), 406–421.

De Valois, K. K., & Switkes, E. (1983). Simultaneous masking interactions between chromatic and luminance gratings. *Journal of the Optical Society of America, 73*, 11–18.

Delahunt, P. B., Webster, M. A., Ma, L., & Werner, J. S. (2004). Long-term renormalization of chromatic mechanisms following cataract surgery. *Visual Neuroscience, 21*(3), 301–307.

Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology, 357*, 241–265.

Ditchburn, R. W., & Ginsborg, B. L. (1952). Vision with a stabilized retinal image. *Nature, 170*(4314), 36–37.

Foster, D. H. (2011). Color constancy. *Vision Research, 51*(7), 674–700.

Freeman, J., & Simoncelli, E. P. (2011). Metamers of the ventral stream. *Nature Neuroscience, 14*(9), 1195–1201.

Georgeson, M. A. (1985). The effect of spatial adaptation on perceived contrast. *Spatial Vision, 1*(2), 103–112.

Goddard, E., Solomon, S., & Clifford, C. (2010). Adaptable mechanisms sensitive to surface color in human vision. *Journal of Vision, 10*(9).

Gollisch, T., & Meister, M. (2010). Eye smarter than scientists believed: Neural computations in circuits of the retina. *Neuron, 65*(2), 150–164.

Hurlbert, A., & Wolf, K. (2004). Color contrast: A contributory mechanism to color constancy. *Progress in Brain Research, 144*, 145–160.

Juricevic, I., & Webster, M. A. (2009). Variations in normal color vision. V. Simulations of adaptation to natural color environments. *Visual Neuroscience, 26*(1), 133–145.

Kay, P., Berlin, B., Maffi, L., Merrifield, W. R., & Cook, R. (2009). *The world color survey*. Stanford: CSLI.

Kay, P., & Regier, T. (2006). Language, thought and color: Recent developments. *Trends in Cognitive Sciences, 10*(2), 51–54.

Kingdom, F. A. (2011). Lightness, brightness and transparency: A quarter century of new ideas, captivating demonstrations and unrelenting controversy. *Vision Research, 51*(7), 652–673.

Kohn, A. (2007). Visual adaptation: Physiology, mechanisms, and functional benefits. *Journal of Neurophysiology, 97*(5), 3155–3164.

Kompaniez, E., Abbey, C. K., Boone, J. M., & Webster, M. A. (2013). Adaptation and visual search in medical images. *Journal of Vision, 13*(9), 1232.

Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience, 10*(6), 779–786.

Krauskopf, J., & Gegenfurtner, K. (1992). Color discrimination and adaptation. *Vision Research, 32*(11), 2165–2175.

Krauskopf, J., Williams, D. R., Mandler, M. B., & Brown, A. M. (1986). Higher order color mechanisms. *Vision Research, 26*(1), 23–32.

Kuehni, R. G. (2004). Variability in unique hue selection: A surprising phenomenon. *Color Research and Application, 29*, 158–162.

Kwon, M., Legge, G. E., Fang, F., Cheong, A. M., & He, S. (2009). Adaptive changes in visual cortex following prolonged contrast reduction. *Journal of Vision, 9*(2) (article 20).

Land, E. H. (1986). Recent advances in Retinex theory. *Vision Research, 26*(1), 7–21.

Laughlin, S. (1981). A simple coding procedure enhances a neuron's information capacity. *Zeitschrift fur Naturforschung C – A Journal of Biosciences, 36*(9–10), 910–912.

Lee, T. W., Wachtler, T., & Sejnowski, T. J. (2002). Color opponency is an efficient representation of spectral properties in natural scenes. *Vision Research, 42*(17), 2095–2103.

Lennie, P. (2003). The cost of cortical computation. *Current Biology, 13*(6), 493–497.

Long, F., Yang, Z., & Purves, D. (2006). Spectral statistics in natural scenes predict hue, saturation, and brightness. *Proceedings of the National Academy of Sciences of the United States of America, 103*(15), 6013–6018.

MacLeod, D. I. A. (2003). Colour discrimination, colour constancy, and natural scene statistics (The Verriest Lecture). In J. D. Mollon, J. Pokorny, & K. Knoblauch (Eds.), *Normal and defective colour vision*. London: Oxford University Press.

McDermott, K., Juricevic, I., Bebis, G., & Webster, M. A. (2008). Adapting images to observers. In B. E. Rogowitz & T. N. Pappas (Eds.), *Human vision and electronic imaging* (pp. V-1–V-10). SPIE, 68060.

McDermott, K. C., Malkoc, G., Mulligan, J. B., & Webster, M. A. (2010). Adaptation and visual salience. *Journal of Vision, 10*(13), 17.

McDermott, K. C., & Webster, M. A. (2012a). The perceptual balance of color. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 29*(2), A108–A117.

McDermott, K. C., & Webster, M. A. (2012b). Uniform color spaces and natural image statistics. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 29*(2), A182–A187.

Miyahara, E., Pokorny, J., Smith, V. C., Baron, R., & Baron, E. (1998). Color vision in two observers with highly biased LWS/MWS cone ratios. *Vision Research, 38*(4), 601–612.

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*(4), 783–792.

Ohzawa, I., Sclar, G., & Freeman, R. D. (1982). Contrast gain control in the cat visual cortex. *Nature, 298*(5871), 266–268.

O'Neil, S. F., & Webster, M. A. (2014). Filling in, filling out, or filtering out: Processes stabilizing color appearance near the center of gaze. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 31*(4), A140–147.

Ranganath, C., & Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience, 4*(3), 193–202.

Regan, B. C., & Mollon, J. D. (1997). The relative salience of the cardinal axes of colour space in normal and anomalous trichromats. In C. R. Cavonius (Ed.). *Colour vision deficiencies* (Vol. XIII, pp. 261–270). Dordrecht: Kluwer.

Reynolds, J. H., & Heeger, D. J. (2009). The normalization model of attention. *Neuron, 61*(2), 168–185.

Richards, W. (1979). Quantifying sensory channels: Generalizing colorimetry to orientation and texture, touch, and tones. *Sensory Processes, 3*(3), 207–229.

Rieke, F., & Rudd, M. E. (2009). The challenges natural images pose for visual adaptation. *Neuron, 64*(5), 605–616.

Rinner, O., & Gegenfurtner, K. R. (2000). Time course of chromatic adaptation for color appearance and discrimination. *Vision Research, 40*(14), 1813–1826.

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*, 2036–2045.

Schwartz, O., Hsu, A., & Dayan, P. (2007). Space and time in visual context. *Nature Reviews Neuroscience, 8*(7), 522–535.

Series, P., Stocker, A. A., & Simoncelli, E. P. (2009). Is the homunculus "aware" of sensory adaptation? *Neural Computation, 21*(12), 3271–3304.

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience, 33*, 89–108.

Shevell, S. K. (2012). The Verriest Lecture: Color lessons from space, time and motion. *Journal of the Optical Society of America A: Optics, Image Science, and Vision, 29*(2), A337–345.

Shevell, S. K., Holliday, I., & Whittle, P. (1992). Two separate neural mechanisms of brightness induction. *Vision Research, 32*(12), 2331–2340.

Shevell, S. K., & Kingdom, F. A. (2008). Color in complex scenes. *Annual Review of Psychology, 59*, 143–166.

Singer, B., & D'Zmura, M. (1994). Color contrast induction. *Vision Research, 34*(23), 3111–3126.

Smithson, H. E. (2005). Sensory, computational and cognitive components of human colour constancy. *Philosophical Transaction of the Royal Society of London B: Biological Sciences, 360*(1458), 1329–1346.

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 Series B – Biological Sciences, 216*(1205), 427–459.

Tkacik, G., Garrigan, P., Ratliff, C., Milcinski, G., Klein, J. M., Seyfarth, L. H., et al. (2011). Natural images from the birthplace of the human eye. *PLoS One, 6*(6), e20409.

von der Twer, T., & MacLeod, D. I. (2001). Optimal nonlinear codes for the perception of natural colours. *Network, 12*(3), 395–407.

Wainwright, M. J. (1999). Visual adaptation as optimal information transmission. *Vision Research, 39*(23), 3960–3974.

Wark, B., Lundstrom, B. N., & Fairhall, A. (2007). Sensory adaptation. *Current Opinion in Neurobiology, 17*(4), 423–429.

Webster, M. A. (1996). Human colour perception and its adaptation. *Network: Computation in Neural Systems, 7*, 587–634.

Webster, M. A. (2011). Adaptation and visual coding. *Journal of Vision, 11*(5), 1–23 (article 3).

Webster, M. A., Halen, K., Meyers, A. J., Winkler, P., & Werner, J. S. (2010). Colour appearance and compensation in the near periphery. *Proceedings of the Royal Society of London Series B – Biological Sciences, 277*(1689), 1817–1825.

Webster, M. A., & Juricevic, I. (2013). Optimizing visual performance by adapting images to observers. *SPIE, 8651*(26), 21–28.

Webster, M. A., Juricevic, I., & McDermott, K. C. (2010). Simulations of adaptation and color appearance in observers with varying spectral sensitivity. *Ophthalmic and Physiological Optics, 30*(5), 602–610.

Webster, M. A., & Kay, P. (2007). Individual and population differences in focal colors. In R. E. MacLaury, G. V. Paramei, & D. Dedrick (Eds.), *Anthropology of color* (pp. 29–53). Amsterdam: John Benjamins.

Webster, M. A., & Leonard, D. (2008). Adaptation and perceptual norms in color vision. *Journal of the Optical Society of America A, 25*(11), 2817–2825.

Webster, M. A., & MacLeod, D. I. A. (2011). Visual adaptation and face perception. *Philosophical Transaction of the Royal Society of London B: Biological Sciences, 366*(1571), 1702–1725.

Webster, M. A., Malkoc, G., Bilson, A. C., & Webster, S. M. (2002). Color contrast and contextual influences on color appearance. *Journal of Vision, 2*(6), 505–519.

Webster, M. A., Mizokami, Y., & Webster, S. M. (2007). Seasonal variations in the color statistics of natural images. *Network, 18*(3), 213–233.

Webster, M. A., & Mollon, J. D. (1994). The influence of contrast adaptation on color appearance. *Vision Research, 34*(15), 1993–2020.

Webster, M. A., & Mollon, J. D. (1995). Colour constancy influenced by contrast adaptation. *Nature, 373*(6516), 694–698.

Webster, M. A., & Mollon, J. D. (1997). Adaptation and the color statistics of natural images. *Vision Research, 37*(23), 3283–3298.

Webster, M. A., Werner, J. S., & Field, D. J. (2005). Adaptation and the phenomenology of perception. In C. Clifford & G. Rhodes (Eds.). *Fitting the mind to the world: Adaptation and aftereffects in high-level vision, advances in visual cognition series* (Vol. 2, pp. 241–277). Oxford: Oxford University Press.

Webster, M. A., Werner, J. S., & Field, D. J. (2005b). Adaptation and the phenomenology of perception. In C. W. G. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world: Adaptation and aftereffects in high level vision*. Oxford: Oxford University Press.

Webster, M. A., & Wilson, J. A. (2000). Interactions between chromatic adaptation and contrast adaptation in color appearance. *Vision Research, 40*(28), 3801–3816.

Werner, J. S. (1996). Visual problems of the retina during ageing: Compensation mechanisms and colour constancy across the life span. *Progress in Retinal and Eye Research, 15*(2), 621–645.

Werner, J. S., & Schefrin, B. E. (1993). Loci of achromatic points throughout the life span. *Journal of the Optical Society of America A, 10*(7), 1509–1516.

Williams, D., Tweten, S., & Sekuler, R. (1991). Using metamers to explore motion perception. *Vision Research, 31*(2), 275–286.

Wilmer, J. B. (2008). How to use individual differences to isolate functional organization, biology, and utility of visual functions; with illustrative proposals for stereopsis. *Spatial Vision, 21*(6), 561–579.

Wissig, S. C., Patterson, C. A., & Kohn, A. (2013). Adaptation improves performance on a visual search task. *Journal of Vision, 13*(2), 6.

Wuerger, S. M. (1996). Color appearance changes resulting from iso-luminant chromatic adaptation. *Vision Research, 36*(19), 3107–3118.