

Adaptation and perceptual norms

Michael A. Webster*, Maiko Yasuda, Sara Haber, Deanne Leonard, Nicole Ballardini
Department of Psychology / 296, University of Nevada-Reno, Reno NV 89557, USA

ABSTRACT

We used adaptation to examine the relationship between perceptual norms – the stimuli observers describe as psychologically neutral, and response norms – the stimulus levels that leave visual sensitivity in a neutral or balanced state. Adapting to stimuli on opposite sides of a neutral point (e.g. redder or greener than white) biases appearance in opposite ways. Thus the adapting stimulus can be titrated to find the unique adapting level that does not bias appearance. We compared these response norms to subjectively defined neutral points both within the same observer (at different retinal eccentricities) and between observers. These comparisons were made for visual judgments of color, image focus, and human faces, stimuli that are very different and may depend on very different levels of processing, yet which share the property that for each there is a well defined and perceptually salient norm. In each case the adaptation aftereffects were consistent with an underlying sensitivity basis for the perceptual norm. Specifically, response norms were similar to and thus covaried with the perceptual norm, and under common adaptation differences between subjectively defined norms were reduced. These results are consistent with models of norm-based codes and suggest that these codes underlie an important link between visual coding and visual experience.

Keywords: Vision, adaptation, norm-based coding, color, blur, face perception

1. INTRODUCTION

Many sensory dimensions are thought to be represented relative to a norm, which itself appears psychologically neutral or unbiased. In vision, classic examples of these norms include “white” in the perception of color, or “static” in the perception of motion. Norms have also been implicated in the representation of higher-level stimulus attributes. Thus current models of face perception assume that faces are represented by their “identity trajectories” relative to a prototype or average face [1, 2]. Such diverse examples suggest that norms may in fact be central to most perceptual judgments.

Models of visual coding often assume that a perceptual norm reflects a unique response state in the underlying neural code. That is, the norm for a particular dimension “looks” special because the neural response to this stimulus is qualitatively distinct from the responses to other stimulus levels. In particular, the norm is assumed to reflect a null point or zero crossing in the channel responses, or alternatively, equal responses across the set of channels. As illustrated in Figure 1a, models of color vision include both types of response norms in the representation of white [3]. At the level of the receptors, white is perceived when the responses are balanced across the three classes of cones, whereas any chromatic sensation instead corresponds to a bias in the cone responses. In this case the norm can be thought of as represented only implicitly in the channel responses, since it depends on the distribution of responses across the set of channels. Yet at the post-receptoral level, white is instead encoded by the null point in opponent channels that receive signals of opposite sign from different cone types and thus respond to different hues with either excitation or inhibition. In such channels the achromatic stimulus is uniquely represented by the absence of a response. Thus in this case the norm is represented explicitly, by the (lack of) activity within a channel.

* mwebster@unr.nevada.edu; phone 775-682-8691; fax 775-784-1126

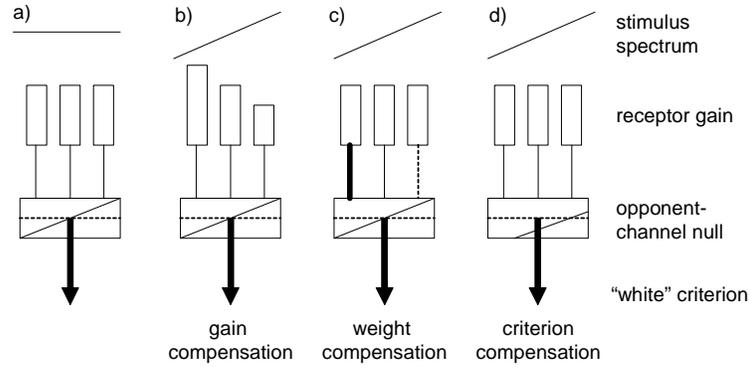


Fig. 1. Norms in color coding. a) The perception of white is assumed to correspond to balanced responses across the three classes of receptors, and to a null point in the response of opponent channels. For a biased spectrum, the responses can be renormalized by changing the gain in the receptors so that their responses are again balanced for the current stimulus (b); changing the weights of the receptor inputs to the opponent channel so that the response to the current stimulus is nulled (c); or changing the criterion, or the relationship between the perceptual norm and the pattern of responses (d).

An alternative to norm-based codes is central tendency models, in which the stimulus is represented by the pattern of activity across multiple channels (Figure 2a) [4]. These models can explain why some dimensions do not have a well defined perceptual norm. For example, in a grating pattern no particular spatial frequency appears special in the way that white does. This is accounted for by assuming that frequency is encoded by the responses across a wide range of channels each narrowly tuned to a different frequency. As a result, any given frequency is represented by the responses across a subset of the channels, and thus no level gives rise to a qualitatively distinct response state. However, note that even for such models a unique norm in the response does occur if a different stimulus is considered. For example, suppose that the stimulus is a broad spectrum of frequencies rather than a single value (Figure 2b) – a stimulus which is much more representative of the natural visual environment [5]. In this case the overall spectrum might be biased toward higher or lower frequencies (so that the image appears blurred or sharpened respectively), and a norm now occurs for a flat spectrum for which the response across the set of channels is equal (e.g. so that the image appears “focused”) [6]. Thus whether or not a response norm might exist depends on both the nature of the stimulus (punctate vs. broadband; Figure 2b) and the nature of the channels (e.g. many narrowly tuned vs few broadly tuned spanning the dimension; Figure 2c).

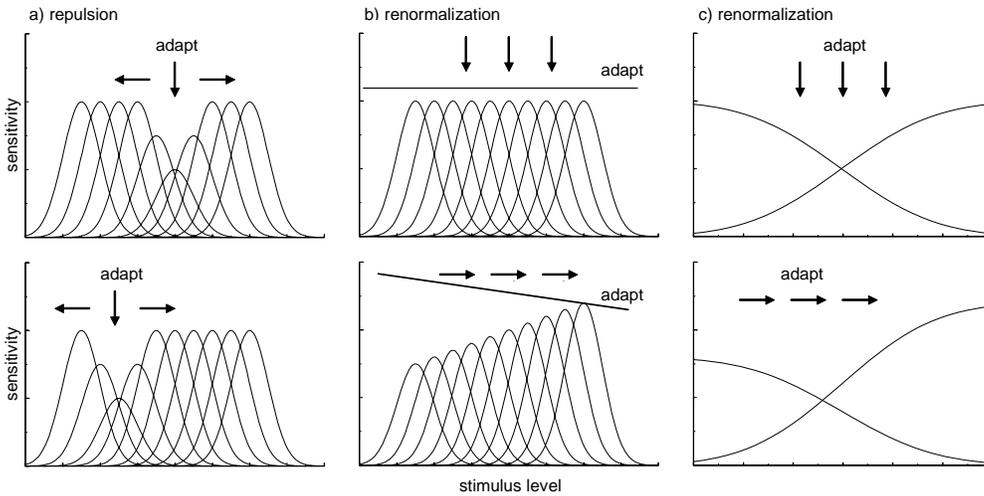


Fig. 2. Central tendency vs. norm based codes. a) multiple channel models assume that the stimulus is represented by activity within a subset of channels tuned to different ranges along the stimulus dimension. This predicts that no stimulus is special, and adaptation leads to a local loss in sensitivity to the adapting level and repulsion away from the adapting level in aftereffects. b) However, the same model leads to a norm-based code and a renormalization in adaptation if the stimulus is broadband or c) if the channels are broadband.

A final point to note is that in natural broadband stimuli the spectrum is typically not flat. For example, in the spatial frequency spectra of natural images the amplitude typically falls in proportion to frequency or as $1/f$ [5, 7, 8]. However, cells in visual cortex may increase their sensitivity roughly in proportion to their preferred frequency (or as $f/1$), so that the response across the set of channels to the average natural spectrum is flat, and thus is at the norm [5]. Indeed, in spatial vision a $1/f$ spectrum “appears” to have equal contrast at all scales, while white noise appears overpowered by the high frequencies. Similarly, in color vision the sensitivities of the cones might be weighted to produce balanced responses to the average color they are exposed to, which will not be spectrally flat (Figure 1B) [9-11]. This illustrates that the norm depends on the match between the stimulus spectrum and the intrinsic neural sensitivity.

While such models continue to dominate accounts of visual coding, the links between perceptual norms and response norms are not in fact well established, and it remains possible that many of the stimuli that appear special do not actually arise from a special neural response. Color vision again provides a telling example. Conventional color-opponent models hold that the unique hues (i.e. pure red, green, blue or yellow) represent the equilibrium point of the red-green and blue-yellow opponent channels [12]. Thus unique yellow or blue is perceived when the red-green channel is nulled. However, the spectral sensitivities of retinal and geniculate cells do not predict the unique hue axes [13, 14]. These cells receive antagonistic inputs from the L and M cones, or signals from the S cones opposed by the L and M cones [14]. The stimuli that isolate the LvsM or SvsLM axes vary perceptually between red vs cyan and purple vs lime, respectively [15-17]. This means that with the possible exception of “red,” the hues that appear unique occur when neither cone-opponent mechanism is nulled, and conversely, the stimuli that do null the neural response do not have any special perceptual salience. This problem is further complicated at the level of cortex, where the early cone-opponent channels are probably elaborated to form multiple color mechanisms each tuned to a different direction in color space [3, 18, 19]. At this level there may be no stimulus that can isolate a single color-opponent channel, and any hue may be represented by multiple-channels similar to models that are envisioned for the encoding dimensions like spatial frequency or orientation [20, 21].

The discrepancies between color appearance and color coding have suggested that the special salience of the unique hues may be tied to special properties of the environment rather than the neural responses. For example, the perceptual blue-yellow axis lies close to the daylight locus and thus may be a learned characteristic of natural illuminants [22, 23]. This learning provides an alternative basis for perceptual norms – that they represent special response criteria independent of visual sensitivity. Indeed, learned criteria could be a common source of norms. They might also explain why individuals with very different visual capacities subjectively describe the world in the same way. For example, the unique hues are largely uncorrelated with visual sensitivity [24-27], and even observers with severe color deficiencies can still name colors in ways that are concordant with color-normal observers [28]. Thus white may be special only because it is the average of the colors we are exposed to. This is illustrated in the final panel of Figure 1 (d), in which the stimulus does not produce equal responses in the cones or a null at an color-opponent site, yet is nonetheless chosen as the achromatic stimulus by the observer.

In this study we directly tested whether perceptual norms are tied to special states in visual sensitivity. To do this we relied on the fact that norms can be strongly biased by varying the observer’s state of adaptation [29]. These adaptation effects can be used to reveal the underlying neutral points in visual sensitivity. For example, after viewing a red field all colors look greener, while adapting to green causes all colors to look redder. There must therefore be an intermediate level that does not bias the observer’s settings. This level uniquely leaves the visual response to color balanced, and thus defines the response norm at the visual sites at which the adaptation is influencing sensitivity. Our aim was to examine whether these response norms are consistent with the observer’s perceptual norm, when they are “not adapted.” In the case of color vision, this meant comparing the white settings under dark adaptation, while no adapting light is present, to the same settings in the presence of different adapting backgrounds. The dark adapted settings measure the subject’s intrinsic white point, while the adapted settings reveal which stimulus level leaves this intrinsic sensitivity unbiased, and we asked whether these two alternative measures of the norm are the same.

These comparisons are more powerful when applied to differences between observers or the same observer under different conditions. Individual differences in perception are pronounced, but the bases for these remain unknown, and are central to the debate over physiological vs. cultural constraints on perception [30]. Figure 3 compares predictions for observers who have the same sensitivity but different criteria (3a) for white or vice versa (3b). That is, one observer

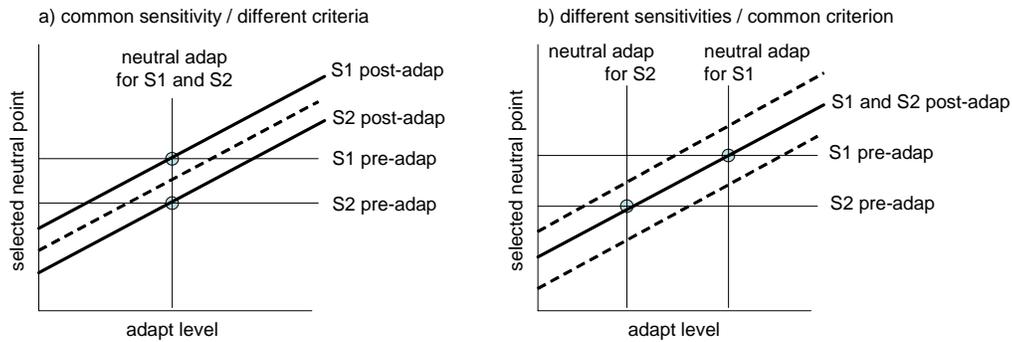


Fig. 3. Predicted effects of adaptation when observers have the same sensitivity but different criteria for a norm, vs. the same criterion but different sensitivities. a) If sensitivity is the same, then the neutral points for both observers will be shifted in the same way, differences between their neutral points will remain before and after adaptation, and the perceptual nulls will not be predictable from the adaptation nulls. b) If the differences instead result from sensitivity differences, then the adaptation nulls will correspond to the subjective nulls, and under adaptation the observers' settings will converge.

might choose a “redder” white point than another either because they have different criteria or different sensitivities. In the former case the neutral stimulus for adaptation will be the same for both observers and thus unrelated to their individual settings, and exposure to this adapting null will not reduce the differences between them (3a). In the latter the neutral points in the adaptation will differ – each observer’s adapting null will be predictable from their perceptual null, and adapting to a common intermediate stimulus will collapse their perceptual nulls to a common point (3b). This leads to two testable predictions for adaptation effects on the perceptual norms across observers: if the appearance differences are due to sensitivity differences (at the level of the adaptation), then 1) the perceptual norms will be correlated with the sensitivity null (as measured by the adaptation) and 2) adaptation to a common stimulus will reduce the variance in perceptual judgments between observers.

Finally, the same logic can be applied to evaluating norms within an observer. For example, sensitivity varies dramatically with retinal location [31], yet perceptually the world itself does not “seem” different in the fovea and periphery, and thus the visual system compensates for the changes in sensitivity with eccentricity to maintain perceptual constancy [32, 33]. The basis for this compensation and the sites at which it might occur again remain largely unknown. In the following experiments, we illustrate three different cases where adaptation can be used to explore these questions and thus help reveal the nature of perceptual norms.

2. METHODS

We measured perceptual norms and adaptation effects for three very different visual judgments: white settings in color vision, focus settings when viewing blurred or sharpened images, and perceived ethnicity when viewing blends of Japanese and Caucasian faces. Details specific to each experiment are described below.

White settings. Chromatic stimuli were displayed on a Sony 20SE monitor controlled by a Cambridge Research Systems Visual Stimulus Generator, which allowed colors on the screen to be controlled with high resolution. The stimulus consisted of a uniform circular 2-deg field on a black background. The field luminance was 25 cd/m². Observers viewed the display binocularly in a dark room and through a hood that shielded extraneous light. The field was either viewed directly or at an eccentricity of 8 deg, in which case a dim gray fixation cross was also displayed. On each run the observer first adapted for 2 min. A test field was displayed for 0.5 sec and then interleaved with brief adaptation top-ups. This sequence continued while the observer adjusted the test field until it appeared achromatic. This was done using a pair of buttons that varied the chromaticity along the two axes of the CIE 1976 u’v’ color space, chosen so that the button presses corresponded to roughly equal perceptual steps. Once the white point was found the response was recorded, and the field was reset to a new random starting chromaticity. A run terminated after 6 white settings. During adaptation subjects viewed either the uniform dark background for their “pre-adapt” settings or the 2 deg field set to the adapting chromaticity. These were chosen to vary chromatic contrast along the LvsM and SvsLM

cardinal axes over a range of ± 80 units in a threshold-scaled version of the MacLeod-Boynton [34] color space centered on the chromaticity of equal energy white. The conversion between our space and the r,b coordinates of MacLeod-Boynton space are given by:

$$\text{LvsM} = 2754 * (r - 0.6657) \quad \text{and} \quad \text{SvsLM} = 4099 * (b - 0.01545)$$

Image focus. Focus settings followed the procedures described in Webster, Georgeson, and Webster [6]. The image was a foliage scene filtered by multiplying the original amplitude spectrum by $f^{-\alpha}$ where f was the spatial frequency. α was varied in steps of 0.01 over a range from -0.5 to +0.5 to create a series of images that varied from moderately blurred to moderately sharpened. After filtering, all images were re-adjusted to have the same rms contrast as the original. The images were shown as a 2-deg field either centrally fixated or centered at an eccentricity of 8 deg. The images were again displayed on a Sony monitor but on a gray background with the same mean luminance as the image and viewed binocularly. The adaptation sequence was similar to the sequence for color. However, in this case subjects made a forced choice response to each presented test image to decide whether it appeared too sharp or too blurred. Subsequent images were varied in a staircase to find the point of subjective focus. Settings were made after adapting to the same image blurred or sharpened over a range from -0.48 to +0.48, with 4 settings made for each adapting level.

Ethnicity judgments. These experiments followed the procedures detailed in Webster et al. [35]. Images were frontal view portraits of Japanese or Caucasian faces cropped to remove the head outline. The images were morphed to form a finely graded series that varied from one ethnicity to the other. Stimuli were shown on the monitor following the same adaptation sequence, and in this case were varied in a staircase while subjects made forced-choice responses to decide whether the displayed image was “Japanese” or “Caucasian.” Settings were repeated before or after adapting to images chosen from different levels of the morph, with 4 settings again made for each adapting level.

3. RESULTS

3.1 White settings

The first example we consider is the perceptual norm of white in color vision. As noted, spectral sensitivity varies across the visual field because of variations in both receptors and post-receptoral processing and in inert screening pigments. In particular, the macular pigment selectively screens light reaching the central macular area of the retina and thus the fovea receives less short wavelength light than the periphery [36]. Despite this, achromatic settings are similar in the fovea and the periphery, indicating that color appearance is compensated for this sensitivity difference [32]. In the same way, color appearance remains stable across the lifespan despite large age-related increases in the density of the lens pigment, which again selectively absorbs short wavelength light and thus strongly biases the spectrum of light reaching the receptors [37]. The basis for this compensation remains unknown, though it has been attributed to more central sites, for example, based on the very slow time course of the renormalization of white settings following cataract surgery [38].

As Figure 1 illustrates, compensation for a spectral sensitivity change could occur at multiple levels. For example, the foveal and peripheral receptors might adjust their intrinsic gains to match the average stimulus they are exposed to (Figure 1a vs. 1b). Alternatively, the receptors might respond the same yet how they are weighted at an opponent site might differ to maintain the null for white (Figure 1a vs. 1c). This has, for example, been postulated to account for long term changes in unique yellow settings after prolonged exposure to a biased color environment [39]. Finally, the constancy for white could reflect an unspecified “criterion” difference in how the same chromatic signals are interpreted in the fovea and periphery (Figure 1a vs 1d).

To distinguish between these, we used chromatic adaptation to steady uniform fields to examine how achromatic settings are biased in the fovea and periphery. An advantage of this is that this adaptation is known to produce largely cone-specific changes in sensitivity and thus can reveal the response norm at a cone specific site [40]. Figure 4a plots the white settings in the LvsM and SvsLM space for a single observer under different states of adaptation. Notably, for this observer the dark adapted settings (large filled symbols) are not completely concordant, yet they are much closer than predicted by the spectral sensitivity change if the fovea and periphery differ in macular screening (shown by the

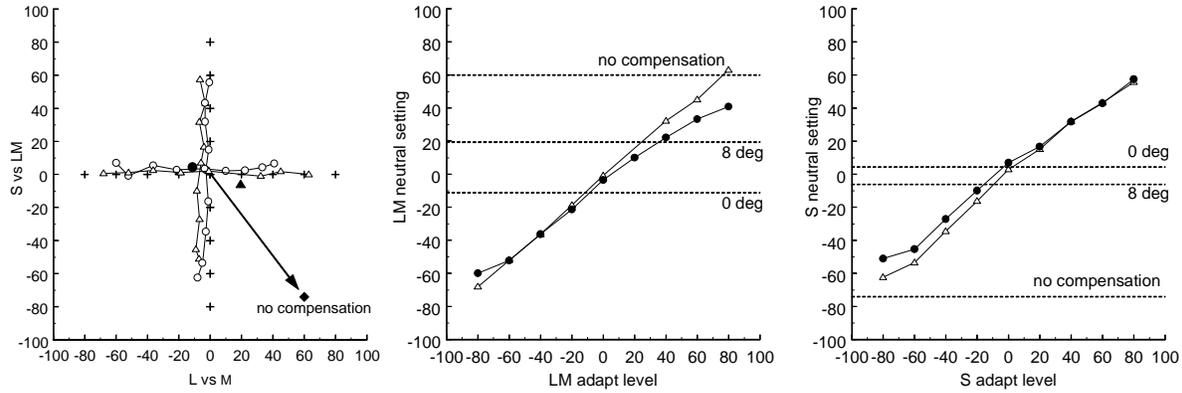


Fig. 4. White settings before or after adaptation to different chromatic levels along the LvsM or SvsLM axes (crosses). a) coordinates of the chosen white points are shown for stimuli presented in the fovea (open circles) or at an eccentricity of 8 deg (open triangles). Large filled triangle and circle show the corresponding achromatic points under dark adaptation. Diamond plots the shift with eccentricity predicted if there is no compensation for macular screening. b) LvsM settings replotted as a function of the LvsM adapting level. Lines show the dark adapted nulls at the two eccentricities or the null predicted for no compensation. c) SvsLM settings plotted as a function of the SvsLM adapting level.

arrow for the standard observer). The unfilled symbols plot the achromatic settings after adapting to chromaticities confined to the LvsM or SvsLM axes. In Figures 4b and 4c these have been replotted to show the change in the LvsM (or SvsLM) coordinate of the white point as a function of the LvsM (or SvsLM) adapting level. Note that in both cases the adapted white settings are similar for the fovea and periphery, and that the neutral adapting level (i.e. the level that does not bias the pre-adapted white settings) remains close to the dark adapted white point. Again, this would not be predicted if the compensation occurred at sites subsequent to the site affected by the adaptation. In that case the adaptation null should be the same when the stimulus reaching the receptors is the same, and again should differ by roughly the difference predicted by macular screening. If the adaptation in this case is indeed cone-specific, then this means that the white settings do correspond to a response norm and that this norm is established at the very initial stages of color coding.

3.2 Focus settings

The same logic can be applied to analyzing the variations in retinal location for other stimulus dimensions. Again, visual acuity falls with increasing eccentricity [31]. Yet these selective losses in contrast sensitivity are not paralleled by subjective impressions of visual focus, because physically focused images in the periphery appear subjectively focused (even though it is obviously not possible to perceive the same level of detail in them). In fact, peripheral targets can appear sharper than blurred foveal targets because of over overcompensation for the sensitivity loss [33].

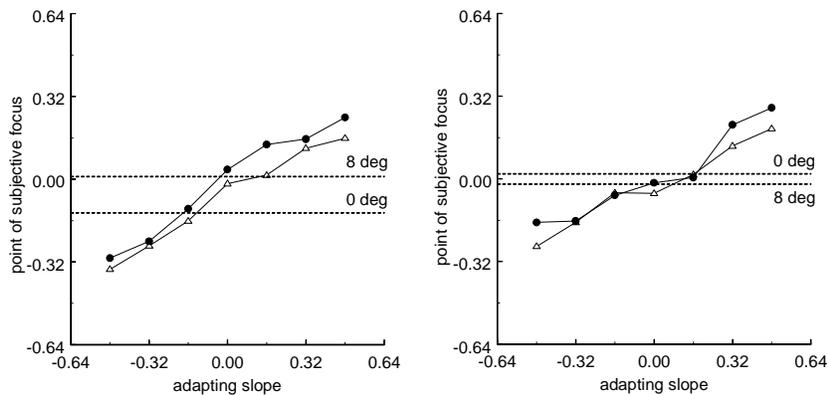


Fig. 5. Focus settings before or after adaptation to blurred or sharpened images presented in the fovea or 8 deg in the periphery. The two panels plot the results for two observers.

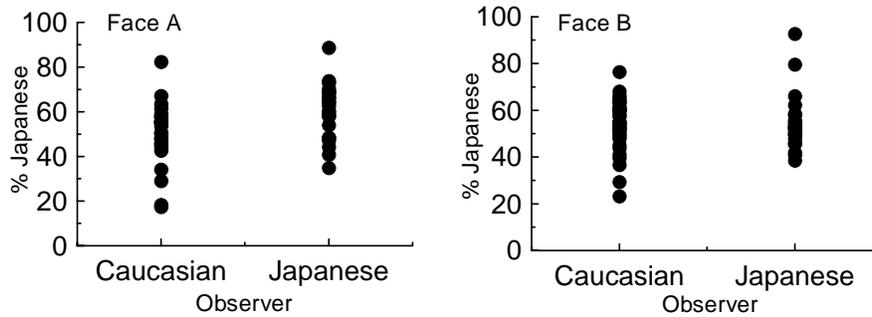


Fig. 6. Subjective category boundaries for ethnicity judged in morphs between a Japanese and Caucasian face for Caucasian or Japanese observers. The two panels plot the settings for a morph between two male (Face A) or two female (Face B) faces. Caucasian observers on average perceive an ambiguous face as more Japanese and vice versa, though this difference was significant only for Face A.

We again asked whether subjective norms for focus could be accounted for by how the spatial mechanisms encoding blur are normalized. In this case subjects adjusted the amplitude spectra of an image so that it appeared focused before or after adapting to blurred or sharpened versions of the image. Results are shown for two observers in Figure 5. Again these results show that adaptation biased the norm in very similar ways at the two eccentricities, and that the adapting blur level that maintained the same pre-adapted focus settings was again close to the pre-adapted focus. Thus like color, this suggests that the norm for color corresponds to the response norm in the spatial channels adapted by the images (though in this case we have not attempted to model the differences that would be predicted if the sensitivity losses are not compensated at the level of the adaptation). However, unlike color, it is probable in this case that the compensation involves a cortical site, since the pattern selective adaptation on which these blur adjustments depend is unlikely to arise in the retina [29].

3.2 Facial categories

In the final example we explored how perceptual norms vary between individuals rather than within. This experiment also differed from the previous two in that we examined the norms for a very high-level visual attribute – the perception of ethnicity in a human face. A final important difference is that this norm reflects a balance between two stimulus categories rather than a centroid for the set of stimuli. That is, while “white” and “focused” may be among the most common stimuli in their class, an image at the boundary between two facial categories may in fact be rare, and the norm instead applies to the balanced response to the two categories. Nevertheless, the same techniques can be used to examine whether a perceptual norm for a facial category can be tied to a response norm in the visual representation of faces. As noted, models of face perception assume that faces are encoded relative to a norm in ways that are closely analogous to color coding [41]. Moreover, the norm for a face can be strongly biased by adaptation in much the same way that adaptation biases the perception of white [42]. These adaptation effects occur for natural variations in individual faces and affect natural judgments about faces, including their gender, ethnicity and identity [35, 43]. For example, in a previous study we found that prior adaptation to a Japanese or Caucasian face induces a large negative aftereffect in the appearance of an ethnically ambiguous face image [35, 43]. That is, after viewing a Japanese face an ambiguous face appears more Caucasian and vice versa. Moreover, we also found that individuals differ in the stimulus levels defining their category boundaries for ethnicity. Intriguingly, Japanese observers are more likely to classify an ambiguous face as Caucasian, while Caucasian observers show on average the opposite bias. This suggests that the observers are more sensitive to how the face differs from their own ethnic category. In the present study we asked whether these differences result from differences in the underlying sensitivity to faces vs. differences in criteria for labeling faces.

Subjects were Caucasian undergraduates at the University of Nevada or Japanese exchange students who were tested within two weeks of their arrival in the US. The observers set the neutral point for morphs between face images of two different Japanese and Caucasian individuals from the Matsumoto and Ekman JACNEUF face set. Pair A were male faces while Pair B were female. Figure 6 plots the neutral points for each face pair for each group of observers. For the male pair (A) the average settings were significantly different, and again in the direction that observers’ boundaries

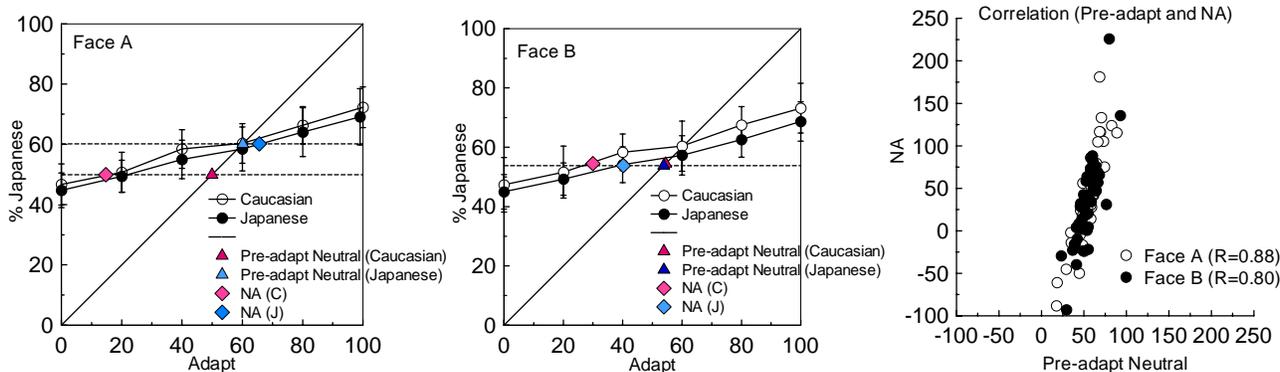


Fig. 7. Shifts in ethnicity boundaries with adaptation. a) Settings in the category boundary are plotted as a function of the level of the adapting morph for Caucasian (open circles) or Japanese (filled circles) observers for face pair A. Dashed lines plot the pre-adapt category boundary. b) Settings for face pair B. c) Scatterplot of the relationship between the pre-adapt ethnicity boundaries and the neutral adapting level for individual observers.

were shifted in the direction of their own ethnicity. However, group differences for the second, female pair (B) were not significant.

Figure 7a and 7b compare the ethnicity boundaries before or after adapting to different levels of the ethnicity morphs. The two curves plot the average settings for each ethnic group. As found for the white and focus settings, the settings under adaptation are very similar and the neutral adaptation levels are close to the level chosen as norms before adaptation. Note that for face pair A, where the groups differ, this means that when adaptation was neutral (on average) for one group, it was biased for the other. For example, if subjects were adapted to the mean neutral point for Caucasians, then on average this did not bias the settings for Caucasian observers, yet did induce an aftereffect (on average) for the Japanese observers and vice versa. (Note that this applies to the average group aftereffect, but not to all individuals since the within-group variance was pronounced.) Thus this suggests that part of the difference in how the faces are classified might in fact be tied to the underlying visual sensitivity of the observers to the face stimuli. This is further seen in Figure 7c which shows that individual differences in the pre-adapt boundaries were correlated with differences in the neutral adapting levels (i.e. the adapt level that did not shift their category boundary). This correlation is again consistent with a sensory difference as the basis for the neutral points. However, the adapt neutral points varied much more than the subjective neutral points, and for some subjects the predicted neutral adapt level fell outside the range spanned by the two original faces. Thus a subject's perceptual neutral point does not completely predict their neutral adapting stimulus. One possible reason for this is that the faces vary along many unknown perceptual dimensions, yet our stimuli and subjects' responses were constrained to a single axis (and one that confounds many dimensions such as ethnicity and identity).

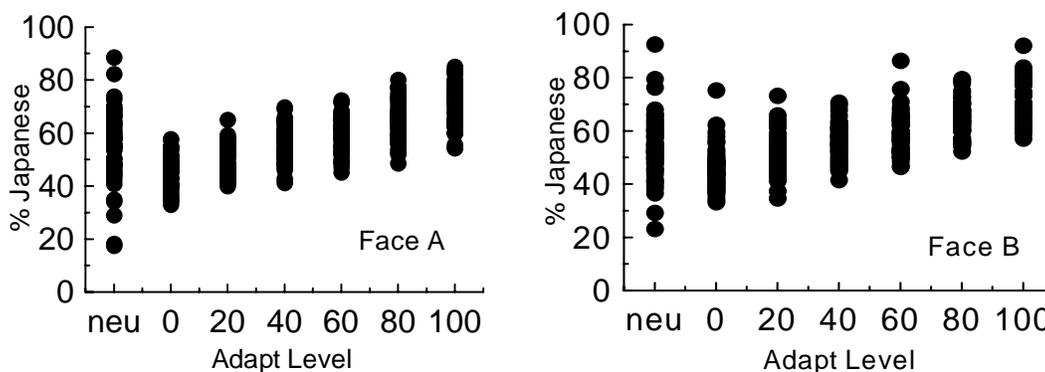


Fig. 7. Comparisons of the range of settings for the ethnicities boundaries before (neu) or after adaptation to different levels of the face morphs. Variance decreased by a factor of 4 for face pair A at all adaptation levels and decreased by a factor of 1.5 to 4 for face pair B.

As we noted above, if the norm is determined by the underlying state of adaptation, then adapting observers to a common stimulus should reduce the differences in their settings (since the external adapting stimulus will drive them all toward a common sensitivity). Figure 8 shows that this was in fact the case for the ethnicity ratings for both face pairs. In particular, for pair A the variance between all observers was a factor of 4 or more lower after adaptation to any of the morph levels than for the pre-adapt settings. These results suggest that at least part of the variation in how observers categorize faces may depend on actual differences in how faces are perceptually encoded, rather than merely differences in criterion.

4. DISCUSSION AND CONCLUSIONS

In this study we have explored the link between visual perception and visual coding by asking whether stimuli that appear psychologically neutral to an observer reflect special neutral response states in visual coding. Each of the domains we tested is consistent with the hypothesis that perceptual norms are determined at least in part by actual response norms. This causal relationship is implicated by the findings that 1) the stimulus that observers perceive as a norm is close to the same stimulus that leaves the underlying state of adaptation in balance; 2) consistent perceptual norms between the fovea and periphery are paralleled by consistent response norms in the adaptation, even though the effective retinal stimulus differs markedly because of eccentricity-dependent filtering effects; 3) individual differences in perceptual norms are correlated with the individual differences in neutral adapting levels; and 4) adaptation to a common stimulus reduces differences in the perceptual neutral points, suggesting that the common effect on sensitivity underlies the convergence in subjective visual judgments. Because we found these results for very different classes of stimuli, it seems likely that many if not most perceptual norms may behave in similar ways and have a similar basis. Moreover, this supports the assumption that norm-based codes represent an important and general strategy involved in both the neural representation and phenomenal experience of visual stimuli.

One advantage of probing norms with adaptation is that adaptation to different stimuli alters sensitivity at different sites in the visual system. For example, while chromatic adaptation to a uniform field primarily affects sensitivity at the level of the receptors or cone-specific pathways, adapting to chromatic contrast instead biases sensitivity at post-receptoral and primarily cortical loci [44]. This makes it possible to probe the response norms at different stages of visual coding and to identify the stages that are likely to define perceptual norms. The fact that cone-specific adaptation adjusts color appearance relative to the same norm that determines the observer's subjective norm for white implies that this norm is established at a very early stage in color coding.

An important function of normalization in perception may be to maintain perceptual constancy despite changes in the viewing context or the observer [45]. For example, maintaining the norm for white or focus may allow stable perceptions of the world despite the pronounced optical and neural changes that occur during development and aging. The same processes may also prove fundamental for providing a form of perceptual constancy between observers. For example, we showed that inter-observer agreement in ethnicity judgments is far greater when the observers are adapted to the same stimulus. It may be this common adaptation to a common world that allows different individuals to have shared perceptual experiences [45].

5. ACKNOWLEDGMENTS

Supported by National Eye Institute grant EY-010834.

REFERENCES

1. Leopold, D.A., et al., *Prototype-referenced shape encoding revealed by high-level aftereffects*. *Nature Neuroscience*, 2001. **4**: p. 89-94.
2. Valentine, T., *A unified account of the effects of distinctiveness, inversion, and race in face recognition*. *Quarterly Journal of Experimental Psychology*, 1991. **43A**: p. 161-204.
3. Webster, M.A., *Human colour perception and its adaptation*. *Network: Computation in Neural Systems*, 1996. **7**: p. 587-634.

4. Braddick, O., F.W. Campbell, and J. Atkinson, *Channels in vision: Basic aspects*, in *Handbook of Sensory Physiology VIII*, R. Held, H.W. Leibowitz, and H. Teuber, Editors. 1978, Springer-Verlag: Berlin. p. 3-38.
5. Field, D.J., *Relations between the statistics of natural images and the response properties of cortical cells*. Journal of the Optical Society of America A, 1987. **4**: p. 2379-2394.
6. Webster, M.A., M.A. Georgeson, and S.M. Webster, *Neural adjustments to image blur*. Nature Neuroscience, 2002. **5**: p. 839-840.
7. Tolhurst, D.J., Y. Tadmor, and T. Chao, *Amplitude spectra of natural images*. Ophthal. Physiol. Opt., 1992. **12**: p. 229-232.
8. van der Schaaf, A. and J.H. van Hateren, *Modelling the power spectra of natural images: Statistics and information*. Vision Research, 1996. **36**: p. 2759-2770.
9. Brown, R.O., *The world is not grey*. Investigative Ophthalmology and Visual Science (Suppl.), 1994. **35**: p. 2165.
10. Judd, D.B., D.L. MacAdam, and G. Wyszecki, *Spectral distribution of typical daylight as a function of correlated color temperature*. Journal of the Optical Society of America, 1964. **54**: p. 1031-1040.
11. Webster, M.A. and J.D. Mollon, *Adaptation and the color statistics of natural images*. Vision Research, 1997. **37**: p. 3283-3298.
12. Hurvich, L.M. and D. Jameson, *An opponent-process theory of color vision*. Psychological Review, 1957. **64**: p. 384-404.
13. De Valois, R.L., I. Abramov, and G.H. Jacobs, *Analysis of response patterns of LGN cells*. Journal of the Optical Society of America, 1966. **56**: p. 966-977.
14. Derrington, A.M., J. Krauskopf, and P. Lennie, *Chromatic mechanisms in lateral geniculate nucleus of macaque*. Journal of Physiology, 1984. **357**: p. 241-265.
15. Krauskopf, J., D.R. Williams, and D.W. Heeley, *Cardinal directions of color space*. Vision Research, 1982. **22**: p. 1123-1131.
16. Webster, M.A., et al., *Variations in normal color vision. A comparison of unique hues in Indian and U.S. observers*. Invest. Ophthal. Visual Sci., 2001. **42**: p. S98.
17. De Valois, R.L., et al., *Hue scaling of isoluminant and cone-specific lights*. Vision Research, 1997. **37**: p. 885-897.
18. Krauskopf, J., et al., *Higher order color mechanisms*. Vision Research, 1986. **26**: p. 23-32.
19. Lennie, P., *Color coding in the cortex*, in *Color Vision: From Genes to Perception*, K.R. Gegenfurtner and L.T. Sharpe, Editors. 1999, Cambridge University Press: Cambridge. p. 235-247.
20. Webster, M.A. and J.D. Mollon, *Changes in colour appearance following post-receptoral adaptation*. Nature, 1991. **349**: p. 235-238.
21. Webster, M.A. and J.D. Mollon, *The influence of contrast adaptation on color appearance*. Vision Research, 1994. **34**: p. 1993-2020.
22. Mollon, J.D., *Color vision*. Annual Review of Psychology, 1982. **33**: p. 41-85.
23. Pokorny, J. and V.C. Smith, *Evaluation of single-pigment shift model of anomalous trichromacy*. Journal of the Optical Society of America, 1977. **67**: p. 1196-1209.
24. Brainard, D.H. and B.A. Wandell, *Analysis of the retinex theory of color vision*. Journal of the Optical Society of America A, 1986. **3**: p. 1651-1661.
25. Miyahara, E., et al., *Color vision in two observers with highly biased LWS/MWS cone ratios*. Vision Research, 1998. **38**: p. 601.
26. Scheffrin, B.E. and J.S. Werner, *Loci of spectral unique hues throughout the life span*. Journal of the Optical Society of America A, 1990. **7**: p. 305-311.
27. Webster, M.A., et al., *Variations in normal color vision. II. Unique hues*. Journal of the Optical Society of America A, 2000. **17**: p. 1545-1555.
28. Nolan, J.N., M.A. Crognale, and M.A. Webster, *Color appearance and central color coding in a case of cone dysfunction syndrome*. International Colour Vision Society Proceedings, 2002.
29. Webster, M.A., *Pattern selective adaptation in color and form perception*, in *The Visual Neurosciences Volume 2*, L.M. Chalupa and J.S. Werner, Editors. 2003, MIT Press: Cambridge. p. 936-947.
30. Berlin, B. and P. Kay, *Basic Color Terms: Their Universality and Evolution*. 1969, Berkeley: University of California Press.
31. De Valois, R.L. and K.K. De Valois, *Spatial Vision*. 1988, Oxford: Oxford University Press.
32. Beer, R.D., et al., *Compensation of white for macular filtering*. Journal of Vision, 2005. **5**: p. 282a.
33. Galvin, S.J., et al., *Sharpness overconstancy in peripheral vision*. Vision Research, 1997. **37**: p. 2035-2039.
34. MacLeod, D.I.A. and R.M. Boynton, *Chromaticity diagram showing cone excitation by stimuli of equal luminance*.

- Journal of the Optical Society America, 1979. **69**: p. 1183-1186.
35. Webster, M.A., et al., *Adaptation to natural facial categories*. Nature, 2004. **428**: p. 558-561.
 36. Snodderly, D.M., J.D. Auran, and F.C. Delori, *The macular pigment. II. Spatial distribution in primate retinas*. Investigative Ophthalmology and Visual Science, 1984. **25**: p. 674-685.
 37. Werner, J.S. and B.E. Scheffrin, *Loci of achromatic points throughout the life span*. Journal of the Optical Society of America A, 1993. **10**: p. 1509-1516.
 38. Delahunt, P.B., et al., *Long-term renormalization of chromatic mechanisms following cataract surgery*. Visual Neuroscience, 2004. **21**: p. 301-307.
 39. Neitz, J., et al., *Color perception is mediated by a plastic neural mechanism that is adjustable in adults*. Neuron, 2002. **35**: p. 783-792.
 40. Stiles, W.S., *Color vision: The approach through increment-threshold sensitivity*. Proceedings of the National Academy of Sciences, 1959. **45**: p. 100-114.
 41. Rhodes, G., et al., *Adaptation and face perception: how aftereffects implicate norm-based coding of faces*, in *Fitting the Mind to the World: Adaptation and Aftereffects in High Level Vision*. 2005, Oxford University Press: Oxford. p. 213-240.
 42. Webster, M.A. and O.H. MacLin, *Figural after-effects in the perception of faces*. Psychonomic Bulletin and Review, 1999. **6**: p. 647-653.
 43. Leopold, D.A., et al., *Prototype-referenced shape encoding revealed by high-level aftereffects*. Nature Neuroscience, 2001. **4**: p. 89-94.
 44. Webster, M.A. and J.A. Wilson, *Interactions between chromatic adaptation and contrast adaptation in color appearance*. Vision Research, 2000. **40**: p. 3801-3816.
 45. Webster, M.A., J.S. Werner, and D.J. Field, *Adaptation and the phenomenology of perception*, in *Fitting the Mind to the World: Adaptation and Aftereffects in High Level Vision*, C.W.G. Clifford and G. Rhodes, Editors. 2005, Oxford University Press: Oxford. p. 241-277.