Response normalization and blur adaptation: Data and multi-scale model

Sarah L. Elliott,
Institute for Mind and Biology, University of Chicago, Chicago, IL, USA

Mark A. Georgeson, and
School of Life and Health Sciences, Aston University, Birmingham, UK

Michael A. Webster
Department of Psychology, University of Nevada, Reno, NV, USA

Abstract

Adapting to blurred or sharpened images alters perceived blur of a focused image (M. A. Webster, M. A. Georgeson, & S. M. Webster, 2002). We asked whether blur adaptation results in (a) renormalization of perceived focus or (b) a repulsion aftereffect. Images were checkerboards or 2-D Gaussian noise, whose amplitude spectra had (log–log) slopes from −2 (strongly blurred) to 0 (strongly sharpened). Observers adjusted the spectral slope of a comparison image to match different test slopes after adaptation to blurred or sharpened images. Results did not show repulsion effects but were consistent with some renormalization. Test blur levels at and near a blurred or sharpened adaptation level were matched by more focused slopes (closer to $1/f$) but with little or no change in appearance after adaptation to focused ($1/f$) images. A model of contrast adaptation and blur coding by multiple-scale spatial filters predicts these blur aftereffects and those of Webster et al. (2002). A key proposal is that observers are pre-adapted to natural spectra, and blurred or sharpened spectra induce changes in the state of adaptation. The model illustrates how norms might be encoded and recalibrated in the visual system even when they are represented only implicitly by the distribution of responses across multiple channels.

Keywords

spatial vision; blur adaptation; response normalization

Introduction

Blur is an intrinsic and important feature of image quality and is an attribute that observers are highly sensitive to (Watt & Morgan, 1983, 1984; Wuerger, Owens, & Westland, 2001). The image formed on the retina is inherently blurred due to aberrations in the cornea and lens of the eye, the finite aperture of the pupil, fluctuations in accommodation, and limited depth of focus. In addition, information from the natural environment may be blurred when factors such as motion or mist are present in the scene. Despite the prominence of such imperfections in the retinal image, most observers do not report the world as appearing out of focus, and even observers with substantial refractive errors or neural acuity deficits may

not normally experience the world as blurry. In fact, human observers are very good at judging whether an image itself is in proper focus (Field & Brady, 1997; Tadmor & Tolhurst, 1994).

The appearance of correct focus might reflect learning the average blur we are exposed to and associating that with the structure of the world. An alternative is that visual coding is adapted to compensate for retinal image blur, in the same way that color appearance is compensated for the spectral biases present in the scene (e.g., because of the illumination) or the eye (e.g., because of filtering by the lens or macular pigment). There is now substantial evidence that the visual system does adapt or adjust to changes in the level of blur. For example, adaptation to optically induced blur has an effect on acuity (George & Rosenfield, 2004; Mon-Williams, Tresilian, Strang, Kochhar, & Wann, 1998; Pesudovs & Brennan, 1993; Rajeev & Metha, 2010) and contrast sensitivity (Mon-Williams et al., 1998; Rajeev & Metha, 2010); and adapting to images with varying levels of blur induces strong biases in the shape of the contrast sensitivity function measured both psychophysically (Webster & Miyahara, 1997; Webster, Mizokami, Svec, & Elliott, 2006) and in single cells in primary visual cortex (Sharpee et al., 2006). Moreover, adaptation to blurred images has a dramatic effect on the appearance of blur (Battaglia, Jacobs, & Aslin, 2004; Elliott, Hardy, Webster, & Werner, 2007; Vera-Diaz, Woods, & Peli, 2010; Webster, Georgeson, & Webster, 2002; Webster et al., 2006). Specifically, after adapting to images that are blurred (or sharpened) by "distorting" the ratio of low to high spatial frequency content, a physically focused image appears too sharp (or blurred), so that the point of best subjective focus is shifted toward the prevailing frequency content of the adapting images.

However, the consequences of these adaptation effects for the perception of image focus remain unclear. If adaptation serves to "discount" blur in the retinal image so that the world appears focused, then the adapting images themselves should appear less blurred and better focused as observers adjust to them. This would result in a "renormalization" of perceived focus so that the adapting stimulus appears more neutral. Alternatively, the adaptation could instead reflect a selective loss in sensitivity to the adapting blur level. This could leave the perceived level of adapting blur unchanged, while inducing a "repulsion" aftereffect in nearby blur levels, because the distribution of responses to different blur levels is biased away from the adapting level (Graham, 1989). For example, images physically sharper (or more blurred) than the adapting level might appear even sharper (or even more blurred). Webster et al. (2002) informally tested these alternatives by asking subjects to rate the perceived focus of images after a period of prolonged exposure and found that the images were subjectively judged as less blurred, consistent with renormalization. However, their experiment left unresolved whether the reported shifts reflected a change in perception or criterion. Moreover, both patterns of adaptation predict similar aftereffects in a physically focused image (i.e., in both cases a blurred or sharp adapter will cause a focused image to appear sharper or blurrier, respectively). Thus, prior measurements of blur adaptation with focused tests cannot discriminate between the models. To accomplish this, in the present study, we used an asymmetric matching task in order to measure how the adaptation altered the appearance of blur over a wide range of test blur levels, for which the two forms of adaptation make different predictions.

The form of the adaptation is important for understanding both the functional implications of the adaptation and the representation of blur in the visual system. Renormalization would imply that the point of perceived focus reflects a perceptual norm that appears neutral and qualitatively distinct from other levels along the stimulus dimension (Webster & Leonard, 2008). This is similar to the special nature of "gray" as a norm in color vision. Norms are typically modeled as a balanced response across a small number of broadly tuned channels (e.g., two channels tuned to blurrier or sharper), or as the null point in a single opponent

channel. In contrast, repulsion aftereffects imply that the stimulus dimension is represented by multiple narrowly tuned mechanisms. In that case, no single level is special and adaptation instead produces a more localized sensitivity loss (though this is again a form of normalization, to the relative sensitivity across the set of channels). Repulsion is in fact characteristic of the aftereffects of size or spatial frequency (Blakemore & Sutton, 1969). Indeed, such aftereffects were central to the development of multi-channel models of spatial vision. How can a norm for blur exist within such models?

One possible answer is that the different behaviors of the broadband vs. narrowband models arise when the stimulus is narrowband (e.g., a single spatial frequency; Figure 1 right). If the stimulus itself is broadband, then a unique norm again exists when the responses are balanced across the set of channels (Figure 1 left). For example, if perceived blur is related to the relative energy at different frequencies, then a multiple channel model might represent a norm when the responses are balanced across the relevant set of channels. An aim of our work was therefore not only to test for the presence of norms in blur adaptation but to assess how these norms could arise given standard multi-channel models of spatial coding.

The representation of blur in the visual system may be carried by both global and local codes (Field & Brady, 1997; Georgeson, May, Freeman, & Hesse, 2007). Natural images— like simple edges—have amplitude spectra characterized by greater energy at lower spatial frequencies and thus follow a roughly inverse relationship between amplitude and frequency. However, the slope of the global amplitude spectrum varies widely across images and thus is itself a poor predictor of image focus (Field & Brady, 1997; Tolhurst, Tadmor, & Chao, 1992). Nevertheless, for any given image, steepening the slope biases the spectrum toward lower frequencies and increases the perception of blur, while shallower slopes conversely increase the relative energy at higher frequencies and perceptually sharpen the image. In the present study, we used these slope changes to manipulate perceived blur and to test how different levels of blur are affected by adaptation. While these variations do not simulate actual optical blur, nor Gaussian blur, they have the advantage that they span both blurred and sharp directions relative to the original image and thus may more directly tap the neural calibrations underlying perceptual judgments of focus.

We also explored adaptation to these spectral slope variations for two classes of images— edges and filtered spatial noise. Blurring or sharpening an image can change many attributes of an image. These changes include variations in the spatial profile of edges and in the apparent texture contrast at different spatial scales. They can also include changes in perceived shape, for example, with astigmatic blur (Anstis, 2002; Sawides et al., 2010). It is not clear which of these attributes might drive the adaptation nor whether they might adapt in similar ways. For example, in simple step edges, the spatial frequency components are all in phase and blurring or sharpening produces a localized luminance change. When an edge is blurred, the luminance change becomes more gradual and the width of the transition is increased. When adapting to blur in edges, observers might be encoding and adapting to the altered luminance profile or the local scale of the edge (Georgeson et al., 2007) rather than the global amplitude spectrum. For noise images, however, the perception of blur or sharpness might be carried by the salience of texture at different scales (Field & Brady, 1997). For example, a blurred noise image appears to have less "speckle," perhaps reflecting a more global coding of spatial scale information.

We measured the aftereffects of adapting to blur for filtered noise and for light–dark checkerboards to try to isolate different potential cues to focus. For both, our results suggest that the adaptation does tend toward renormalization of perceived focus, and we show that the data can be explained rather accurately by a model of the adaptive changes in the

contrast response of spatial filter mechanisms that are pre-calibrated to the average amplitude spectrum of focused images.

Methods

Observers

Four female observers participated in the main experiment (mean age = 29.3, range 25–36 years), all having normal or corrected-to-normal visual acuity. Observers were recruited from the University of Nevada, Reno.

Apparatus and stimuli

Stimuli were presented on a gamma-corrected 50-cm Sony Triniton color monitor driven by a PC with 8-bit color resolution. The original images consisted of 2 sets corresponding to: (1) 50 grayscale images of noise filtered to have an amplitude spectrum of $1/f$ (Webster & Miyahara, 1997), and (2) two grayscale images of a checkerboard pattern, one a contrast reversal of the other. The noise images were constructed by first creating white noise with pixel intensities chosen from random normal deviates and then filtering to a $1/f$ spectrum. The images had a fixed root mean square contrast of 0.35, chosen to avoid significant truncation (<0.5%) of the pixel intensities before or after filtering. The images were adjusted to a mean gray level of 128, corresponding to a luminance of 15 cd/m^2, and were presented on a uniform gray background with the same luminance.

All images contained 256×256 pixels, subtended 4° in a field corresponding to 256×256 pixels on the monitor, and had a maximum spatial frequency of 32 c/deg. To blur or sharpen, each image was filtered by multiplying the original amplitude spectrum by f^α, where f is spatial frequency in cycles per degree, and α controls the magnitude of change (Field & Brady, 1997; Knill, Field, & Kersten, 1990; Tadmor & Tolhurst, 1994; Webster et al., 2002). The absolute (log–log) spectral slope was thus equal to $\alpha - 1$. For the test stimuli, α was varied from −0.75 (i.e., absolute slope = −1.75, appearing strongly blurred) to +0.75 (i.e., absolute slope = −0.25, appearing strongly sharpened) in steps of 0.25. Examples from the 2 filtered image sets are shown in Figure 2. For the comparison images, spectral slopes were varied in steps of 0.01 to allow finely graded adjustments in blur magnitude.

Procedure (Experiment 1)

Observers viewed the screen binocularly from a distance of 112 cm, and their basic task was to adjust the spectral slope of a comparison image in a trial-by-trial staircase procedure, to match the perceived blur or sharpness of a test image that was presented after exposure to an adapting image. There were three adaptation conditions: (a) adaptation to the original images ($1/f$ spectrum; $\alpha = 0$); (b) adaptation to blurred images ($\alpha = -0.5$); and (c) adaptation to sharpened images ($\alpha = +0.5$). For each adapting condition, there were 7 test conditions defined by 7 α levels from $\alpha = -0.75$ to $+0.75$ (see Figure 2). Slope matches for each test level were estimated 3 times for a total of 126 trials (7 test levels $\times$ 3 adaptation levels $\times$ 2 image sets $\times$ 3 repeats) for each observer. Sessions began with an adaptation period of 120 s, which displayed either a random sequence of the noise images or a counterphasing sequence of the checkerboard images, at only one of the adaptation levels: (1) blurred ($\alpha = -0.5$), (2) original ($\alpha = 0$), or (3) sharpened ($\alpha = +0.5$) in different sessions. The inner edges of the adaptation images were located 1° to the left of a 0.34° fixation cross, accompanied by a uniform gray field (gray level 128) to the right of fixation. For adaptation to noise, the random sequence involved resampling a noise image from the set every 0.25 s, to homogenize local light adaptation and minimize afterimages. There was a 0.25-s blank gray screen between adaptation and test images.

The 0.5-s test phase displayed one of the 7 test α levels to the left of fixation (same location as the adaptation sequence) and a variable comparison image to the right of fixation (previously a uniform field). Using a two-alternative forced-choice staircase method, the observer's task was to adjust the α of the comparison image on the right to match the perceived α of the test image on the left by pressing buttons to indicate whether the comparison image appeared sharper or more blurred. During a single session, two staircases with a 1-up, 1-down procedure were randomly interleaved, and each adjustment changed the spectral slope of the right-hand image by an α of 0.01. Each staircase terminated after 8 reversals. The test phases were interleaved with 6-s periods of top-up adaptation. The mean α from the last 6 of 8 reversals was used as the estimate of the observer's perceived match to the test image.

Control experiment (Experiment 2)

A potential problem with the adaptation configuration just described is that observers might be adapting differently to the contrast of the adapting image (to the left of fixation) and the zero-contrast uniform field (to the right of fixation). A difference in perceived contrast might affect the judgments of blur (e.g., because lower contrast images might appear sharper). To reduce this possibility, the second experiment used a random sequence of control adapting images ($1/f$ spectrum; $\alpha = 0$) to the right of fixation, instead of a uniform field. One observer (YM) was recruited from the first experiment and the second was author MW. Both observers had normal or corrected-to-normal visual acuity. Procedures were otherwise the same as the first experiment, except that each session was repeated twice using only the noise image set to give a total of 42 settings (7 test levels × 3 adaptation levels × 1 image set × 2 repeats).

Results

Experiment 1: Blur matches

The two leftmost panels in Figure 3 show the perceived α matches following adaptation for the two image sets, averaged over the 4 observers (checkerboard: top left; spatial noise: bottom left). Significant adaptation effects can be seen for matches made during the blurred ($\alpha = -0.5$) and sharpened ($\alpha = +0.5$) adaptation conditions ($F_{(2,6)} = 21.01$, $p < 0.01$ ANOVA, main effect of adaptation condition), irrespective of the image set used for adaptation. In the blurred adaptation condition, perceived α matches are seen to shift toward a higher α value, indicating that the test images appeared sharper after blurred adaptation than after focused or $1/f$ adaptation ($\alpha = 0$). The opposite can be seen for the sharpened adaptation condition; perceived α matches shifted to lower α values, showing that the test images appeared more blurred.

The shift in perceived α matches following adaptation was not significantly different for the two image sets ($F_{(1,3)} = 2$, NS), and results were qualitatively similar for the 4 observers. This point is illustrated for two observers tested with the checkerboards (Figure 3, top middle and right panels) and spatial noise (bottom middle and right panels).

The shifts in the α matches as a function of the test α level are inconsistent with a simple repulsion model. As noted, this idea supposes that adaptation will not alter the perceived level of the adapting stimulus, and that test levels higher or lower than the adapting level will appear shifted in opposite ways. (The predictions for repulsion are illustrated later, in Figure 9.) Instead, the data show that both the blurred and the sharpened adapters appeared more focused after adaptation. These shifts are highlighted in Figure 4, which shows, for either the blurred or sharpened adapter, the difference between mean α matches and a veridical (i.e., physical) match. Adapting to blurred images ($\alpha = -0.5$; red circles) made

most test images seem sharper but not to equal extents: the sharpening effect was less strong for test images that were already sharpened (α = +0.25 to +0.75). Similarly, adapting to sharpened images (α = +0.5; blue squares) made sharpened test images seem more blurred but with less effect on blurred test images. This effect of the test α level on the shift away from a physical match following adaptation to blurred or sharpened images was significant ($F_{(12,36)}$ = 3.4, $p < 0.01$ ANOVA). Adaptation had stronger effects at or near the adapting level than on test levels that were far removed. The aftereffects thus exhibit a renormalization, but not one that is uniform across all blur levels.

If the data always exhibited "repulsion," then $1/f$ adapting images (α = 0) should make sharpened test images seem even sharper, and blurred test images more blurred, but this was not so. Mean α matches (absolute values) did not deviate significantly from the physical α level of the test images after adaptation to a sequence of $1/f$ images ($t_{(1,12)}$ = 2.18, $p = 0.8$). The lack of significant aftereffects from $1/f$ adapting images is consistent with the idea of renormalization, since this adapting level is already at the norm and so should not induce a recalibration. We explore this in the model described below.

Experiment 2: Controlling for adaptation to contrast

As noted, a possible complication in Experiment 1 is that the aftereffects might be influenced by changes in apparent contrast as well as apparent Fourier spectral slope. The comparison stimuli were shown on a previously uniform (zero-contrast) field while the test stimuli were shown on a field preceded by the adapting images, and so, because of contrast adaptation, the test images probably had a lower perceived contrast, and this might affect perceived blur (May & Georgeson, 2007). To control for this, the experiment was repeated with $1/f$ noise images presented on the comparison side during adaptation while the adapting noise images were shown on the test side.

Similar aftereffects were found for these conditions (Figure 5). Adapting to $1/f$ (α = 0), images again produced little shift in perceived α matches, but this must now be expected (by symmetry) because the adapting images were the same on both sides of the display. Adapting to sharpened images caused the test images to appear more blurred (Figure 5B, blue squares), and adapting to blurred images produced a sharpened aftereffect (Figure 5B, red circles). Like Experiment 1, these effects were not uniform across the test levels but were instead stronger for test levels near the adapting level. Thus, this experiment confirmed that the effects of adaptation persisted when potential effects of contrast differences during adaptation were controlled.

Summary of results

The main features of our experimental results are: (i) adapting to a $1/f$ or in-focus image (α = 0) produced little or no systematic change in perceived blur of any test image; (ii) adapting to a sharpened image (α = +0.5) made $1/f$ or sharpened images seem more blurred but had much less impact on physically blurred ones; (iii) adapting to a blurred image (α = −0.5) did the reverse, making $1/f$ or blurred images seem sharper but with a much smaller effect on physically sharpened ones. In the next section, we develop a quantitative, multi-scale model of contrast adaptation and blur coding to account for the observed aftereffects.

Modeling the blur aftereffects

After observers have adapted to a moderate- or high-contrast grating for a few minutes, the contrast threshold for gratings of similar orientation and spatial frequency is raised (Blakemore & Campbell, 1969; Pantle & Sekuler, 1968) and perceived contrast of such gratings is lowered (Blakemore, Muncey, & Ridley, 1973; Georgeson, 1985). Cells in the primary visual cortex show a reduced response after grating adaptation, and the change in

contrast-response function can be characterized as a multiplicative change in contrast gain, or in response gain, or more often both (Albrecht, Farrar, & Hamilton, 1984; Dean, 1983; Ohzawa, Sclar, & Freeman, 1982). Some cortical cells show a shift in their peak or preferred spatial frequency or orientation after adaptation to off-peak stimuli, implying that for these cells adaptation involves more than a simple reduction in responsiveness. Such shifts in tuning seem to occur in complex cells rather than simple cells (Movshon & Lennie, 1979; Muller, Metha, Krauskopf, & Lennie, 1999).

In psychophysics, several descriptive models of contrast adaptation are possible (Foley & Chen, 1997). We asked whether the observed changes in perceived blur and sharpness might be explained by the known psychophysical properties of contrast adaptation, as applied to a multi-channel system for encoding blur. For ease of exposition with minimal complexity, we begin with a simple, 1-parameter, descriptive model of contrast adaptation. Georgeson (1985) found that the reductions in perceived contrast of sine-wave gratings after adaptation to similar gratings could not be described by a multiplicative change in contrast gain but could be described by a simple subtractive rule: perceived (matched) contrast of a grating was fairly well predicted by subtracting about one-third of the adapting contrast from the test contrast. Note that this descriptive rule is expressed in terms of physical contrasts, not the responses of visual mechanisms.

Multi-channel model of blur adaptation

Here we are dealing with complex images that presumably stimulate many filters at different spatial scales and orientations, so we elaborated the subtractive rule by supposing that it operates at the level of individual filters. A sketch of the main components of the model is shown in Figure 6, elaborated below. In brief, we propose that adaptation alters the responses of individual filters, and that blur (or sharpness) in these tasks is determined by whether responses rise (or fall) as filter scale increases. If responses rise (or fall) more steeply with increasing scale after adaptation, then perceived blur (or sharpness) is increased. Crucially, however, we find that to account adequately for the results we must suppose that observers are pre-adapted to a focused world before they begin the experiment. We shall propose that blurred or sharpened adapting images whose spectra differ from this *norm* temporarily change the state of adaptation, but focused images or blank images do not.

The model filters had odd-symmetric receptive fields that were first directional derivatives of a 2-D Gaussian function $G(x, y; s)$ (sometimes called "edge detectors"; see Georgeson et al., 2007, their Figure 1A) at 4 orientations (0, 45, 90, and 135 deg from vertical) and 7 scales s (from 1 to 8 pixels in 0.5 octave steps). An expression for the vertically oriented receptive field of scale s is

$$\frac{\partial}{\partial x}\{G(x, y; s)\} = -x \cdot g(s)e^{-x^2/2s^2} e^{-y^2/2s^2}.$$

(1)

For the 4-deg image size used, optimal spatial frequencies of these filters ranged from 10 c/deg ($s = 1$ pixel) to 1.25 c/deg ($s = 8$ pixels). (Note that filter orientations 180, 225, 270, and 315 would be redundant, since those filters differ only in sign from the first four, and we used an unsigned (r.m.s.) response from each filter.) Spatial frequency bandwidth at the optimal orientation is the same for all filters (2.6 octaves full-width at half-height). At small scales, the small receptive fields (RFs) contain just a few significant pixels, and the way in which they are sampled becomes important. We used the optimal method for discretely sampled Gaussians described by Lindeberg (1994). Receptive field amplitudes were scaled by a factor $g(s)$ chosen so that all filters had the same amplitude in the 2-D Fourier domain,

and this meant that the variance or "energy" of the response to 2-D noise with a $1/f$ spectrum was the same for all filters (Field, 1987; Field & Brady, 1997). This can be thought of as a long-term calibration of the filters for natural images, whose *average* amplitude spectrum is very close to $1/f$ [the weighted mean spectral slope from 11 studies, 1176 images, was −1.08 (Billock, 2000)]. The factor $g(s)$ was not varied in the simulations described here.

The images used in a given experiment were convolved with each of the 28 filters, and each filter's response was reduced to a single number r by computing the standard deviation of the filtered image values over all pixel positions, omitting a 16-pixel-wide border to minimize edge truncation artifacts. The value of r can be regarded as a measure of the spatial contrast (or more strictly, the *amplitude* of variation) in the spectral band that is "seen" by a given filter. The model thus adopts a global, rather than local, approach to the encoding of image blur (Field & Brady, 1997). For each test blur, the responses (r) were averaged by taking the r.m.s. value over different test exemplars (32 different noise samples for Experiments 1 and 2; four different test images [face, leaves, meadow, checks] for the experiment of Webster et al., 2002). Since r is computed separately at each scale and orientation, we can define r(test) as the 4×7 *array* of responses across orientations and scales, representing the average pattern of filter responses to a set of test images that had a specific spectral slope. Similarly, r(adapt) is the pattern of responses to a specified adapting slope.

Simple fatigue model

The subtractive model can now be defined as

$$R(\text{test, adapt})=r(\text{test}) - k \cdot r(\text{adapt}). \tag{2}$$

$R(\cdot)$ is the 4×7 array of filter responses to the *test* pattern, as modified by the effects of the *adapt* pattern. Small or negative values of R (below 0.1) were set to an assumed baseline value of 0.1. Equation 2 is an example of a "fatigue model," because each filter's loss of response to the *test* image is a simple function of how much that filter responded previously to the *adapt* image(s). The free parameter k represents the strength of adaptation, and we assumed it to be the same for all filters. For 3 c/deg sine-wave gratings, k was about 0.3 (Georgeson, 1985). When the adapter is blank, r(blank_adapt) = 0, and so R(test, blank_adapt) = r(test).

To predict judgments of blur or sharpness, we adopted an approach similar to that of Field and Brady (1997), by pooling the responses $R(\cdot)$ over the 4 filter orientations and then computing the slope of the (log–log) response over scales. The average (r.m.s.) value of R was calculated over filter orientations, leaving just a single response value $R_{av}(s)$ at each filter scale s. This reduction to a single number, pooled over space and orientation, is not intended as a general model for spatial vision but seems appropriate here where the image filtering is global and isotropic, and the observer makes a simple binary choice between two images. The values of $R_{av}(s)$ were fitted with a power function ($R_{av} = a \cdot s^b$) by finding the least squared error, and the exponent b was taken as the code for perceived blur. Roughly speaking, b estimates blur from the (log) ratio of activity in large- and small-scale filters.

Figure 7A shows how R_{av} varies with filter scale s when nothing is subtracted (blank adaptation). For a $1/f$ image ($\alpha = 0$; filled circles) R_{av} is, by design, constant across scale and so b is close to 0. As spectral blur increases ($\alpha < 0$; open symbols), responses decrease at small scales and slope b becomes increasingly positive. The reverse is true with spectral sharpening ($\alpha > 0$; filled symbols): slope b becomes increasingly negative. Slope b is a

therefore a valid metric for this type of blur, because it varies monotonically with changes in the Fourier spectral slope of the test image. Figures 7B–7D show that this monotonic relation holds true after subtractive adaptation to different spectral slopes, but the response slopes *b* for individual test images are altered by adaptation. Compared with 1/*f* adaptation (Figure 7B), the fitted lines tend to rotate toward positive slopes—meaning greater blur—after sharpened adaptation (Figure 7C) and toward negative slopes—meaning greater sharpness—after blurred adaptation (Figure 7C).

We assume that the observer's decisions about blur (spectral slope) are based on *b*. In the present experiments, images viewed after different adapting conditions should appear to match in blur when they yield the same value of *b*. Thus, the model can generate predictions about the observed blur matches. Figure 8 plots the response slopes *b* as a function of spectral slope offset *α* for each of the 4 adapting conditions of Figure 7, with adapting strength *k* = 0.4. Consider the blurred adapter (red, filled circles): each test slope yields a certain response code *b*, and so on this model the predicted blur match is obtained by reading off the curve that represents the comparison condition, to find the comparison *α* that yields the same response code *b*. For Experiment 1, the relevant curve is for blank adaptation (open diamonds). In this way, using cubic spline interpolation in *Matlab*, the model made predictions that could be compared directly with the experimental data. For Experiment 2, the comparison side was adapted to 1/*f* images, so the predicted blur matches were interpolated from the green curve (adapt *α* = 0) instead.

Failure of the simple fatigue model

Figure 9 shows that the simple fatigue model fails. For Experiment 1, it predicts a strong "repulsion" effect for all three adapters, centered on the point at which the adapt and test blurs are equal (marked by dashed circles in Figure 9A). That is, test images sharper than the adapter should be matched by increasingly sharpened images, and test images more blurred than the adapter should be matched by even more blurred comparison images. The fatigue model predicted a pattern of repulsion that was quite unlike the observed data—with trends in entirely the wrong direction (Figure 9B). Interestingly, if we had used only focused (1/*f*) test images (test *α* = 0 in Figure 9B), we might wrongly conclude that the fatigue model was working quite well. The use of a wide range of test blurs was crucial in showing that the fatigue model fails miserably.

Blur adaptation as adjustment of an internal adaptive norm

In light of this failure, we introduce an extension to the subtractive model based on two key ideas about the normative nature of adaptation that prove to be more successful, and perhaps more interesting, than the simple fatigue model. The two key ideas are: (i) the visual system is *pre-adapted* to natural images (with an average 1/*f* spectrum) before the experiment begins, and (ii) the state of adaptation is altered by exposure to adapting images but is unchanged by exposure to a blank field. This last point amounts to assuming that the pre-adapted state persists in the absence of any evidence to the contrary. Such "storage" of visual adaptation has been observed in a variety of circumstances (e.g., McCollough, 1965; Thompson & Wright, 1994; van de Grind, van der Smagt, & Verstraten, 2004; Wohlgemuth, 1911).

An ensemble of images with an average 1/*f* spectrum may be called the *norm*. Thus, in the format of Equation 2, we have for the *pre-adapted* state:

$$R(\text{test, pre} - \text{adapt}) = r(\text{test}) - k \cdot r(\text{norm}). \tag{3}$$

In practice, the response array r(norm) was calculated as the set of filter responses acquired for $1/f$ test images in the simulated experiment, but we might think of it more generally as the stored, long-term average response array r to natural images. Note that this term *pre-adapts* the system to $1/f$ images, in the same subtractive way as before.

To summarize the norm-based model, after experimental exposure to adapting images, Equation 2 applies; before experimental adaptation, or after blank adaptation, Equation 3 applies; after adaptation to $1/f$ images (equivalent to the norm), the two equations are equivalent. Unlike the first model, the net response $R(\cdot)$ to any test image is now the same for a blank adapter and for a $1/f$ adapter, and both are the same as in the pre-adapted state. Equations 2 and 3 entail a stable perceptual response to $1/f$ images but an adaptive change in response to sharpened or blurred ones. Once the system is pre-adapted to the norm, further adaptation to either $1/f$ images, or blanks, causes no change in the system's response to a test image: $R = r$(test) $- k \cdot r$(norm). However, adaptation to a sharp or blurred adapter does change the state of adaptation and so changes the response to a test image: $R = r$(test) $- k \cdot r$(adapt); the adapter behaves (temporarily) like a new norm.

Predictions of the two models

In both the fatigue and norm-based models, the net outcome $R(\cdot)$ is used to make judgments, via a representation of spectral slope (the blur code b). However, importantly, the two models can make different predictions. When the adapter is blank, as it was on the comparison side in Experiment 1, the norm-based model gives R(test, blank_adapt) = r(test) $- k \cdot r$(norm), while the fatigue model—with no pre-adaptation—gives R(test, blank_adapt) = r(test). These are clearly different and lead to different codes b (compare the green triangles and open diamonds in Figure 8). Thus, we can distinguish the two models when a blank adapter is involved in the experimental task (provided k is not close to 0).

Unlike the fatigue model, the norm-based model fitted the noise-image data of Experiment 1 very closely (Figures 10A and 10D), with only a single free parameter ($k = 0.3$). It captured very well the variable degree of induced blurring for different test images after sharpened adaptation ($\alpha = 0.5$, blue curves), including the slight crossover to sharpening for the most blurred test image (test $\alpha = -0.75$). The complementary effects of adapting to blurred images, including a similar crossover, were also well described by the model (red curves), except for one data point, that may well be an outlier since no such deviation was seen in Experiment 2.

For Experiment 2, where a $1/f$ adapter was shown on the comparison side, the pattern of aftereffects was similar to Experiment 1. Both versions of the model now made the same predictions (Figures 10B and 10E), because, as discussed above, no blank adapter was involved. The aftereffects were a little larger than in Experiment 1 and were well fitted by a small increase in adaptation strength ($k = 0.4$ instead of 0.3). The assumption that observers are pre-adapted to the world (average $1/f$ spectrum) is critical in accounting for these results, because the same model without this assumption failed resoundingly on Experiment 1.

So far, we have discussed the models as applied to noise images. Simulations were also run for the checkerboard images of Experiment 1, as shown in Figures 10C and 10F. The aftereffects were smaller here, and the best fit was obtained with a smaller adaptation strength ($k = 0.15$). Bearing in mind these weaker effects, the norm-based model fit fairly well, except for the most extreme test images ($\alpha = \pm 0.75$). It may be that for images that have distinct, localized features (edges in this case), observers tend to shift from the global to a more local code for blur judgments, but we did not attempt to model this. Importantly, the fatigue model without pre-adaptation again predicted repulsion and was a very poor fit for this condition (not shown) as it was for the noise images (Figure 9B).

Judgments of "best focus" after adaptation

Webster et al. (2002) had observers adjust the spectral slope of natural images until they appeared "in focus," after adaptation to images with various modified spectral slopes. Adaptation to blurred images made in-focus ones seem too sharp, and vice versa. The shifts in focus judgment (Figure 11) suggest partial renormalization but fall short of full normalization (oblique dashed line in Figure 11).

Presumably, observers adjust the spectral slope until it meets some internal, stored standard that represents "in focus." Here the model's estimates of best focus after adaptation were obtained by finding which test stimulus level (α) gave blur code b equal to a stored standard value b_0 given by in-focus images after neutral (either blank or focused) adaptation. This last assumption sidesteps the question of how the visual system knows which images are in focus but enables us to test different models of adaptation. For this task, the norm and no-norm versions of the model have the same behavior because no blank images are used in the experiment, and the standard b_0 is the same for both models.

The good fit of the model to Webster et al.'s (2002) data is shown in Figure 11. Here the adaptation strength $k = 0.5$ was a little higher than in Experiments 1 and 2 ($k = 0.3, 0.4$). For comparison, the dashed curve in Figure 11 shows the somewhat smaller shifts in focus judgments that would be predicted if $k = 0.35$ (mean of Experiments 1 and 2). This difference might simply reflect inter-subject differences in adaptation strength (Vera-Diaz et al., 2010) or differences in the images used (natural images vs. noise). However, another possibility is that part of the effect in Webster et al.'s data reflects higher level criterion shifts (that could be modeled by a shift in the value of the internal standard b_0) in addition to the sensory adaptation that alters blur code b for test images. Matching tasks, like Experiments 1 and 2 here, are likely to be immune from general criterion shifts since the observer is asked to judge one image against another, rather than judge them against an internal criterion. Put simply, if two images appear similar, they should do so whether we judge them to be in focus or not. Thus, the present results are important in showing that the blur aftereffects are not mainly due to criterion shifts. The relatively small difference between dashed and solid curves in Figure 11 could be due to criterion shift, or other factors mentioned above.

In Appendix A, we consider divisive gain controls as alternatives to subtractive adaptation. In brief, we find that a model based on the contrast gain control (CGC) equation of Foley and Chen (1997) can serve as a suitable substitute for subtractive adaptation, but pre-adaptation remains necessary to fit the data of Experiment 1. We have focused primarily on the subtractive model not because it is technically superior, but because of its simplicity. The present data sets are too sparse to allow the various parameters of the CGC model to be reliably or directly estimated.

In the Supplementary material, we develop a dynamic form of the model that gives further insight. It shows how pre-adaptation, and the changing adaptation state, can be driven continuously by the history of stimulation, in a way that is exactly consistent with the static (steady-state) model described here. The dynamic version shows how a key assumption— preservation of the adaptation state during blank intervals—could be produced by a simple storage mechanism that puts that state on "hold" when the input contrast falls to 0, just as the capacitor holds the voltage level in a simple RC filter circuit when the voltage source is switched off.

Summary of the modeling

The norm-based model and fatigue model are closely related and based on the same subtractive principle, but only the norm-based model gives a satisfactory account of all three

experiments (Experiments 1 and 2, and the results of Webster et al., 2002). This model implies that the visual system is pre-adapted to natural images (average $1/f$ spectrum), that adapting to blank or $1/f$ images produces no change in the state of adaptation, and that adapting to sharpened or blurred images does change the state of adaptation and leads to changes in perceived blur or sharpness across a wide range of test images.

General discussion

We found that adaptation to blurred or sharpened images tends to make the adapting images themselves appear better focused (closer to $1/f$) and thus to partially renormalize the subjective point of focus relative to the current adapting level. We showed that these response changes can be closely accounted for by a model based on simple assumptions about contrast adaptation and blur coding by multiple channels in the visual system. Adaptation within these channels acts to reduce imbalance in the distribution of responses to the ambient blur level and thus tends to renormalize the neural code for blur. Our results thus suggest that the perception of image focus, like many other perceptual dimensions, is represented as a norm in visual coding.

Subjective focus as a perceptual and neural norm

Phenomenologically, the point of subjective focus behaves like a norm in having a special and neutral appearance relative to other (blurred or sharpened) stimulus levels; and here we have shown that it also has a special status like other norms in terms of adaptation (in that blurred or sharpened images bias the appearance of focused images but not vice versa), similar to the asymmetries seen with color or face adaptation (Webster & MacLin, 1999). Note that from this perspective the point of best focus is at the "center" of the perceptual continuum. This differs from the perspective suggested by considering only optical factors —where images can become too blurred but never too sharp. Accordingly, most studies of blur have concentrated only on the low-frequency (blurred) side of the representation. However, the neural response can be imbalanced in either direction. It would be instructive to revisit many of the measurements that have characterized blur perception and discrimination to examine performance for stimuli that are instead over-sharpened. This might give a better understanding of the neural encoding of image focus and might also reveal neural responses that are specifically associated only with increasing the low-frequency bias and thus potentially with optical sources of blur.

Norms and multi-scale representations

Typically, norms are assumed to reflect a balance of responses across two broadly tuned channels or to be directly encoded as the null point within an opponent mechanism (Webster & MacLeod, in press). Spatial frequency coding differs in this regard because of the wealth of evidence for multiple narrowly tuned channels representing different spatial scales (reviewed by De Valois & De Valois, 1980; Graham, 1989). In the case of blur, the normalization behavior may arise not from the channel structure alone, but from the fact that the stimulus is broadband. Thus, the norm is again consistent with the simple and general assumption that it reflects balanced or unbiased responses across the set of channels, even if this balance reflects responses in many channels that more finely sample the stimulus dimension. We have shown that the response changes induced by the adaptation can be closely accounted for by a multi-scale model of spatial coding. The main assumptions of the model included (i) spatial filters at multiple scales, (ii) response reduction from adaptation within each filter, and (iii) the encoding of (global) blur from the relative activity across filter scales.

Partial normalization

We have seen that the effect of adaptation on blur did not conform precisely to any of the patterns we had envisaged. Neither repulsion (Figure 9), nor complete normalization (Figure 11), nor uniform normalization across all test blurs (Figure A1) were adequate descriptions of the observed changes in blur matching across the wide range of test blurs used. Instead, our analysis suggests that prolonged adaptation to natural images (pre-adaptation to the *norm*) tends to persist and to *amplify* the blur response to temporary deviations from the normal state. Figure 8 (green triangles) reveals that the blur codes after adaptation to $1/f$ are twice as large ($b \sim = -2\alpha$) as they are without adaptation (open diamonds: $b \sim = -\alpha$). This would make short-term changes in image focus or image quality over time more salient than without adaptation, and that could be useful in both perception and accommodation control. However, if the blurred or sharpened state persists, the system readapts and that salience declines. The adapters come to look more focused, and this can be seen as a form of normalization, but we know rather little about its time course.

Limitations

One objection to the pre-adaptation concept is that, if we are already adapted to in-focus natural images, then further adaptation to such images in the laboratory should have no additional effect. That is what we have assumed here, but Webster and Miyahara (1997) found that exposure to a sequence of unrelated natural images, with a new image every 300 ms, substantially reduced contrast sensitivity and perceived contrast of sine gratings at relatively low spatial frequencies (≤4 c/deg) but not at higher frequencies. Similar findings were reported by Bex, Solomon, and Dakin (2009). It is not easy to dismiss this problem, but its restriction to low spatial frequencies suggests that temporal factors may be important. The rapid sequence of abruptly changing images may have had a spatiotemporal power spectrum with greater energy at high temporal and low spatial frequencies than in natural viewing, where many saccades are small. If so, this may have increased the adaptation level of low SF filters that have transient responses and a greater preference for higher temporal frequencies. A possible resolution of this problem is that the loss of sensitivity to low frequencies observed in these cases might be outside the range of higher frequencies that could be more critical for judging perceived blur and sharpness. A further possibility is that adaptation effects for broadband stimuli, of the kind we have used here, cannot be fully predicted by supposing independent response changes at different spatial frequencies. This was suggested by the finding that adaptation to square waves produced little threshold elevation at the higher harmonics even though adaptation to these harmonics presented alone did reduce sensitivity (Nachmias, Sansbury, Vassilev, & Weber, 1973; Tolhurst, 1972). It is also suggested by the fact that adapting to square wave or $1/f$ patterns does not bias perceived focus but does selectively alter contrast sensitivity.

A final limitation is that the our model was designed only to account for the attribute of global blur and may not predict how adaptation adjusts in other ways to the blur in images. For example, adapting to a single blurred edge has been found to show a repulsion aftereffect, rather than normalization (Georgeson, 2001). Perceived blur of a test edge (assessed by adjusting blur of a comparison edge at an unadapted location) was veridical for test edges whose blur was the same as the adapter but was shifted away from the adapting level for test edges that were more or less blurred, much like the shifts in perceived spatial frequency that follow adaptation to gratings (Blakemore & Sutton, 1969). Thus, local and global blur coding might operate in different ways.

Blur constancy

Renormalization is a common (though not universal) consequence of adaptation across many perceptual dimensions, from color coding to face perception (Webster & Leonard,

2008; Webster & MacLeod, in press). Why are such adjustments so common? In color vision, adapting to the ambient spectrum plays an important role in contributing to color constancy, discounting variations in the stimulus (e.g., the current illumination) to maintain color appearance (e.g., for the same surface; Brainard & Wandell, 1992). These adjustments may be equally important for discounting variations in the observer, for example, over time as visual sensitivity changes with development or aging (Werner & Schefrin, 1993) and over space as spectral sensitivity changes with retinal location (Webster & Leonard, 2008). Similarly, an important functional consequence of renormalization in blur adaptation may be to maintain stable perception of image focus by removing variations caused by the environment or the observer.

For example, the finding that adaptation promotes constancy by causing blurred or sharpened images to appear better focused has important clinical implications for understanding the consequences of refractive errors and their corrections. If visual coding did not adjust to the observer's optical imperfections, then there would be a perpetual mismatch between perception and the world, degrading both subjective image quality and visual performance (even though the full perceptual consequences of this mismatch would not be revealed by standard measures of visual acuity). Moreover, if coding could not readjust after optical correction, then observers might experience the world as unnaturally and uncomfortably sharpened (even if their acuity is necessarily improved). An example of this mismatch has been reported in a congenital cataract patient who did not undergo surgery until middle age and continued to perceive the world as too sharp even months after the correction (Fine, Smallman, Doyle, & MacLeod, 2002). Most observers instead require brief periods to acclimate to a refractive correction, and this adjustment may depend fundamentally on the ability of the visual system to renormalize spatial coding through adaptation.

Acknowledgments

References

Albrecht DG, Farrar SB, Hamilton DB. Spatial contrast adaptation characteristics of neurones recorded in the cat's visual cortex. The Journal of Physiology. 1984; 347:713–739. [PubMed: 6707974]

Anstis SM. Was El Greco astigmatic? Leonardo. 2002; 35:208.

Battaglia PW, Jacobs RA, Aslin RN. Depth-dependent blur adaptation. Vision Research. 2004; 44:113–117. [PubMed: 14637361]

Bex, PJ.; Solomon, SG.; Dakin, SC. Contrast sensitivity in natural scenes depends on edge as well as spatial frequency structure; Journal of Vision. 2009. p. 1p. 1-19.http://www.journalofvision.org/content/9/10/1

Billock VA. Neural acclimation to 1/f spatial frequency spectra in natural images transduced by the human visual system. Physica D. 2000; 137:379–391.

Blakemore C, Campbell FW. On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images. The Journal of Physiology. 1969; 203:237–260. [PubMed: 5821879]

Blakemore C, Muncey JP, Ridley RM. Stimulus specificity in the human visual system. Vision Research. 1973; 13:1915–1931. [PubMed: 4746989]

Blakemore C, Sutton P. Size adaptation: A new aftereffect. Science. 1969; 166:245–247. [PubMed: 5809598]

Brainard DH, Wandell BA. Asymmetric color matching: How color appearance depends on the illuminant. Journal of the Optical Society of America A. 1992; 9:1433–1448.

Dean AF. Adaptation-induced alteration of the relation between response amplitude and contrast in cat striate cortical neurons. Vision Research. 1983; 23:249–256. [PubMed: 6868400]

De Valois RL, De Valois KK. Spatial vision. Annual Review of Psychology. 1980; 31:309–341.

Elliott, SL.; Hardy, JL.; Webster, MA.; Werner, JS. Aging and blur adaptation; Journal of Vision. 2007. p. 8p. 1-9.http://www.journalofvision.org/content/7/6/8

Field DJ. Relations between the statistics of natural images and the response properties of cortical cells. Journal of the Optical Society of America A. 1987; 4:2379–2394.

Field DJ, Brady N. Visual sensitivity, blur and the sources of variability in the amplitude spectra of natural scenes. Vision Research. 1997; 37:3367–3383. [PubMed: 9425550]

Fine I, Smallman HS, Doyle P, MacLeod DI. Visual function before and after the removal of bilateral congenital cataracts in adulthood. Vision Research. 2002; 42:191–210. [PubMed: 11809473]

Foley JM, Chen CC. Analysis of the effect of pattern adaptation on pattern pedestal effects: A two-process model. Vision Research. 1997; 37:2779–2788. [PubMed: 9373676]

George S, Rosenfield M. Blur adaptation and myopia. Optometry & Vision Science. 2004; 81:543–547. [PubMed: 15252354]

Georgeson MA. The effect of spatial adaptation on perceived contrast. Spatial Vision. 1985; 1:103–112. [PubMed: 3940052]

Georgeson, MA. Seeing edge blur: Receptive fields as multi-scale neural templates [Abstract]; Journal of Vision. 2001. p. 438p. 438ahttp://www.journalofvision.org/content/1/3/438

Georgeson, MA.; May, KA.; Freeman, TC.; Hesse, GS. From filters to features: Scale-space analysis of edge and blur coding in human vision; Journal of Vision. 2007. p. 7p. 1-21.http://www.journalofvision.org/content/7/13/7

Graham, NV. Visual pattern analyzers. Oxford, UK: Oxford University Press; 1989.

Knill DC, Field D, Kersten D. Human discrimination of fractal images. Journal of the Optical Society of America A. 1990; 7:1113–1123.

Lindeberg, T. Scale-space theory in computer vision. Dordrecht, Netherlands: Kluwer; 1994.

May KA, Georgeson MA. Blurred edges look faint, and faint edges look sharp: The effect of a gradient threshold in a multi-scale edge coding model. Vision Research. 2007; 47:1705–1720. [PubMed: 17442367]

McCollough C. Color Adaptation of edge detectors in the human visual system. Science. 1965; 149:1115–1116. [PubMed: 17737844]

Mon-Williams M, Tresilian JR, Strang NC, Kochhar P, Wann JP. Improving vision: Neural compensation for optical defocus. Proceedings of the Royal Society B: Biological Sciences. 1998; 265:71–77.

Movshon JA, Lennie P. Pattern-selective adaptation in visual cortical neurons. Nature. 1979; 278:850–852. [PubMed: 440411]

Muller JR, Metha AB, Krauskopf J, Lennie P. Rapid adaptation in visual cortex to the structure or images. Science. 1999; 285:1405–1408. [PubMed: 10464100]

Nachmias J, Sansbury R, Vassilev A, Weber A. Adaptation to square-wave gratings—In search of the illusive third harmonic. Vision Research. 1973; 13:1335–1342. [PubMed: 4722805]

Ohzawa I, Sclar G, Freeman RD. Contrast gain control in the cat visual cortex. Nature. 1982; 298:266–268. [PubMed: 7088176]

Pantle A, Sekuler R. Size-detecting mechanisms in human vision. Science. 1968; 162:1146–1148. [PubMed: 5698858]

Pesudovs K, Brennan NA. Decreased uncorrected vision after a period of distance fixation with spectacle wear. Optometry & Vision Science. 1993; 70:528–531. [PubMed: 8355963]

Rajeev N, Metha A. Enhanced contrast sensitivity confirms active compensation in blur adaptation. Investigative Ophthalmology and Visual Science. 2010; 51:1242–1246. [PubMed: 19797213]

Sawides, L.; Marcos, S.; Ravikumar, S.; Thibos, L.; Bradley, A.; Webster, MA. Adaptation to astigmatic blur; Journal of Vision. 2010. p. 22p. 1-15.http://www.journalofvision.org/content/10/12/22

Sharpee TO, Sugihara H, Kurgansky AV, Rebrik SP, Stryker MP, Miller KD. Adaptive filtering enhances information transmission in visual cortex. Nature. 2006; 439:936–942. [PubMed: 16495990]

Tadmor Y, Tolhurst DJ. Discrimination of changes in the second-order statistics of natural and synthetic images. Vision Research. 1994; 34:541–554. [PubMed: 8303837]

Thompson P, Wright J. The role of intervening patterns in the storage of the movement after-effect. Perception. 1994; 23:1233–1240. [PubMed: 7899039]

Tolhurst DJ. Adaptation to square-wave gratings: Inhibition between spatial frequency channels in the human visual system. The Journal of Physiology. 1972; 226:231–248. [PubMed: 4343317]

Tolhurst DJ, Tadmor Y, Chao T. Amplitude spectra of natural images. Ophthalmic and Physiological Optics. 1992; 12:229–232. [PubMed: 1408179]

van de Grind WA, van der Smagt MJ, Verstraten FAJ. Storage for free: A surprising property of a simple gain-control model of motion aftereffects. Vision Research. 2004; 44:2269–2284. [PubMed: 15208013]

Vera-Diaz FA, Woods RL, Peli E. Shape and individual variability of the blur adaptation curve. Vision Research. 2010; 50:1452–1461. [PubMed: 20417657]

Watt RJ, Morgan MJ. The recognition and representation of edge blur: Evidence for spatial primitives in human vision. Vision Research. 1983; 23:1465–1477. [PubMed: 6666047]

Watt RJ, Morgan MJ. Spatial filters and the localization of luminance changes in human vision. Vision Research. 1984; 24:1387–1397. [PubMed: 6523759]

Webster MA, Georgeson MA, Webster SM. Neural adjustments to image blur. Nature Neuroscience. 2002; 5:839–840.

Webster MA, Leonard D. Adaptation and perceptual norms in color vision. Journal of the Optical Society of America A, Optics, Image Science, and Vision. 2008; 25:2817–2825.

Webster MA, MacLeod DIA. Visual adaptation and face perception. Philosophical Transactions of the Royal Society. in press.

Webster MA, MacLin OH. Figural aftereffects in the perception of faces. Psychonomic Bulletin & Review. 1999; 6:647–653. [PubMed: 10682208]

Webster MA, Miyahara E. Contrast adaptation and the spatial structure of natural images. Journal of the Optical Society of America A, Optics, Image Science, and Vision. 1997; 14:2355–2366.

Webster MA, Mizokami Y, Svec LA, Elliott SL. Neural adjustments to chromatic blur. Spatial Vision. 2006; 19:111–132. [PubMed: 16862835]

Werner JS, Schefrin BE. Loci of achromatic points throughout the life span. Journal of the Optical Society of America A. 1993; 10:1509–1516.

Wohlgemuth A. On the aftereffect of seen movement. British Journal of Psychology, Monograph Supplement. 1911; 1:1–117.

Wuerger SM, Owens H, Westland S. Blur tolerance for luminance and chromatic stimuli. Journal of the Optical Society of America A, Optics, Image Science, and Vision. 2001; 18:1231–1239.

Appendix A

Other models of adaptation

Subtractive adaptation is not the only model of contrast adaptation. Its merits are simplicity (just one free parameter) and linearity (easy to implement; easy to think about). It fits contrast-matching data tolerably well (Georgeson, 1985) but is not suitable for describing contrast discrimination. We therefore explored several other functional forms for contrast adaptation, while leaving the model filters and blur computation unchanged. Using the same notation as Equation 2, a simple divisive gain control for each spatial channel has the following form:

$$R(\text{test, adapt}) = r(\text{test})/\{1 + m \cdot r(\text{adapt})\}. \tag{A1}$$

This model has pure multiplicative (divisive) scaling of test responses driven by the stored response to the adapter. This is analogous to Von Kries scaling in color vision, and the influence of the adapter increases with m. Thus, we found that with $m = 2$, the predictions for Webster et al.'s (2002) experiment on focus judgment (Figure 11) were very close to the dashed line marked "full normalization," but with m reduced to 0.13 the predictions (not shown) were close to the experimental data—partial normalization. This success is not general, however, because on the more wide-ranging test conditions of Experiments 1 and 2, this gain control model fared less well. It predicted an almost parallel shift in matches (Figure A1; partial normalization again), unlike the test-blur-dependent shifts observed.

A more elaborate contrast gain control (CGC) model in the style of Foley and Chen (1997) was more successful:

$$R(\text{test, adapt}) = \frac{r(\text{test})^p}{z^q + r(\text{test})^q + k \cdot r(\text{adapt})^q}.$$

(A2)

When $p - q < 1$, Equation A2 incorporates a compressive contrast response to the test image (a standard feature in models of contrast discrimination), along with further suppression by the adapter. This model was used to fit the adapted contrast-matching data of Georgeson (1985) using stimulus contrast in place of filter response r and a very satisfying fit was obtained with $p = 3$, $q = 2.6$, $z = 0.5$, and $k = 0.16$. Then, with these values of p, q, and z, Equation A2 was used to generate predictions for Experiments 1 and 2. Adaptation strength k was adjusted to obtain a very satisfactory fit for Experiment 2 shown in Figure A2, with $k = 0.2$, and for Experiment 1 (not shown) with $k = 0.1$. The same model gave a good fit to Webster et al.'s (2002) data, with $k = 0.26$; the plot (not shown) was almost identical to Figure 11 (solid curve).

The subtractive and CGC models differ mathematically and do not produce the same form of contrast response either before or after adaptation. However, there is a clear commonality that explains why both are satisfactory in the present context: after adaptation, the reduction in log response level (R) is much larger at low test responses (r) than when r is high. It is this trait that leads to differential changes in log response of small- and large-scale channels in response to blurred (or sharpened) images and, hence, to changes in the response slope b (Figures 7 and 8) that is used to encode blur. Because of response compression ($p - q < 1$), the absolute slopes b are much smaller for the CGC model, but with suitable k the changes in blur b—when mapped back to the equivalent slopes α—emerged as similar for the two models.

Thus, the CGC model (Equation A2) seems to be a suitable substitute for the assumption of subtractive adaptation. It has wider applications in the psychophysics of contrast discrimination and masking but at the expense of greater complexity and less transparency. It is important to note that the idea of pre-adaptation that persists during blank periods is just as vital here as before. When pre-adaptation was removed, the predictions for Experiment 1 exhibited "repulsion" for all three adapting conditions and were just as bad as for the subtractive model shown in Figure 9.
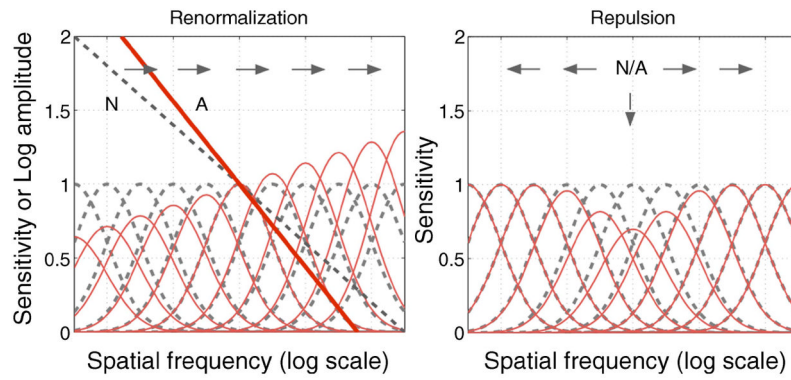
Figure 1.
Some ideas about spatial channels, adaptation and blur. Filter sensitivities (dashed curves) might normally be scaled to give equal responses to the average ($1/f$) spectrum of natural images (dashed line in left panel for the norm (N)). Left: adapting to a steeper, blurred spectrum (solid line, A) alters the relative sensitivities across all channels (red curves) so that responses would be renormalized for the new adapting level. This might cause all slopes (blur levels) to appear shallower (sharper; arrows). Right: basis for the repulsion effect. Adaptation of the same tuned mechanisms to a narrowband stimulus (N/A) instead locally depresses sensitivity to the adapting stimulus. This would bias the distribution of responses away from the adapting level toward either lower or higher perceived frequency (arrows).
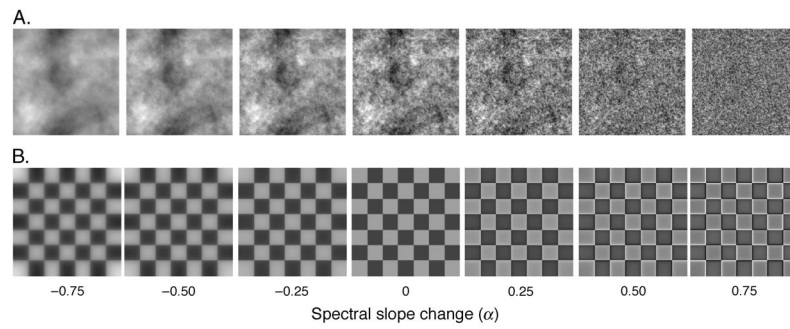
Figure 2.
Example of the 2 image sets and 7 filtered α levels. α values of −0.5 (blurred), 0, and 0.5 (sharpened) were used for adaptation. All 7 filtered α levels were used as test images. (A) Spatial noise images. (B) Checkerboard images.
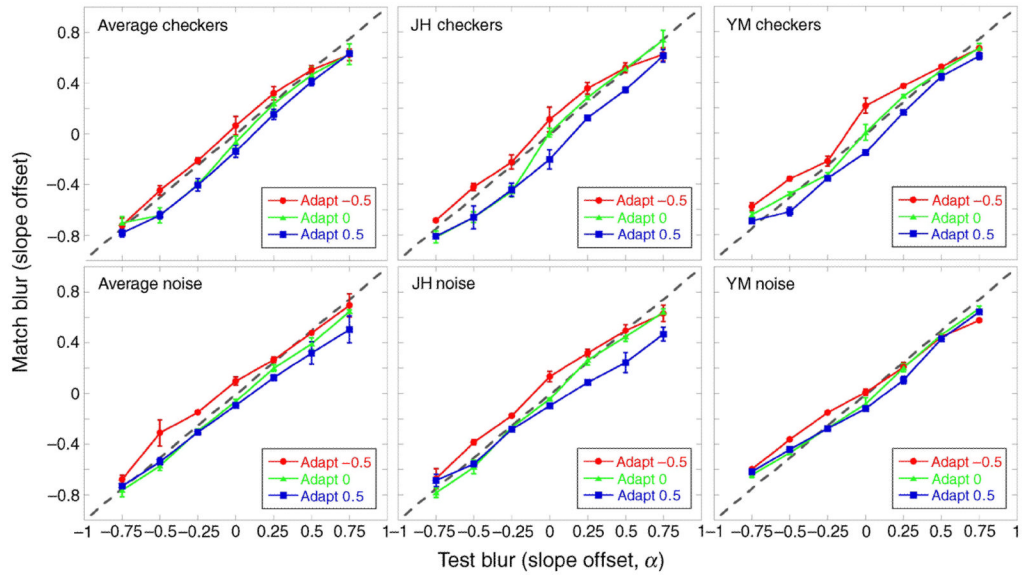
Figure 3.
Experiment 1. Perceived α match after adaptation to blurred (red circles, $\alpha = -0.5$), focused (green triangles, $\alpha = 0$), or sharpened (blue squares, $\alpha = +0.5$) images. The dotted gray line denotes a veridical match. (Top) Checkerboard images. (Bottom) Spatial noise images. (Left) Mean of the 4 observers. (Middle and Right) Two individual observers, JH and YM, respectively. Error bars are ± 1 *SEM*.

Figure 4.
Experiment 1. Mean shifts in perceived α matches after adaptation to blurred (red circles, α = −0.5) and sharpened (blue squares, α = +0.5) images expressed as the difference between observed and veridical matches. The dashed gray line denotes a veridical match. Means for the checkerboards are shown in the left panel, and means for spatial noise are shown in the right panel. Error bars are ±1 *SEM*.
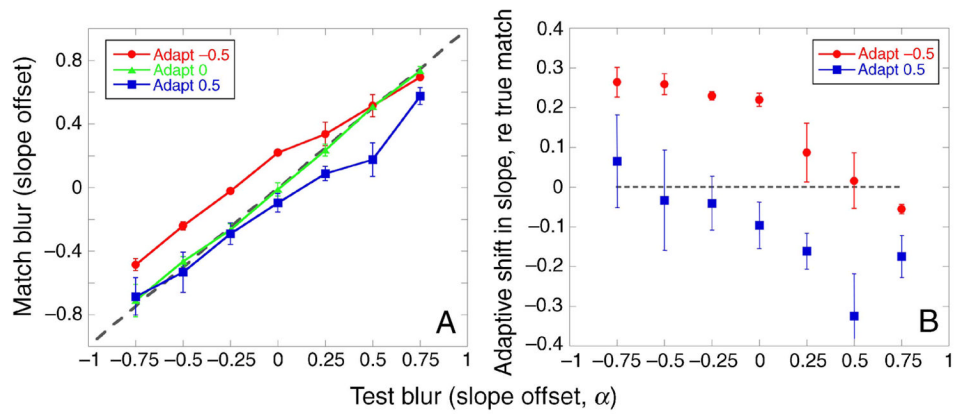
Figure 5.
Experiment 2. (A) Perceived matches following adaptation to blurred (red circles, $\alpha = -0.5$), $1/f$ (green triangles, $\alpha = 0$), or sharpened (blue squares, $\alpha = +0.5$) noise images on the left, and $1/f$ noise images on the right, averaged for the two observers. (B) Mean shifts in perceived matches after adaptation to blurred images (red circles, $\alpha = -0.5$) and sharpened images (blue squares, $\alpha = 0.5$) expressed as the difference between observed and veridical matches. The dotted gray line denotes a veridical match in both panels. Error bars are ± 1 *SEM*.

Figure 6.
A sketch of the multi-scale model for blur adaptation. There were 7 filter scales, at 4 orientations. Strength of adaptation is controlled by factor *k*. Blur is coded from the way the adapted response *R* rises or falls across scales. See text for details.
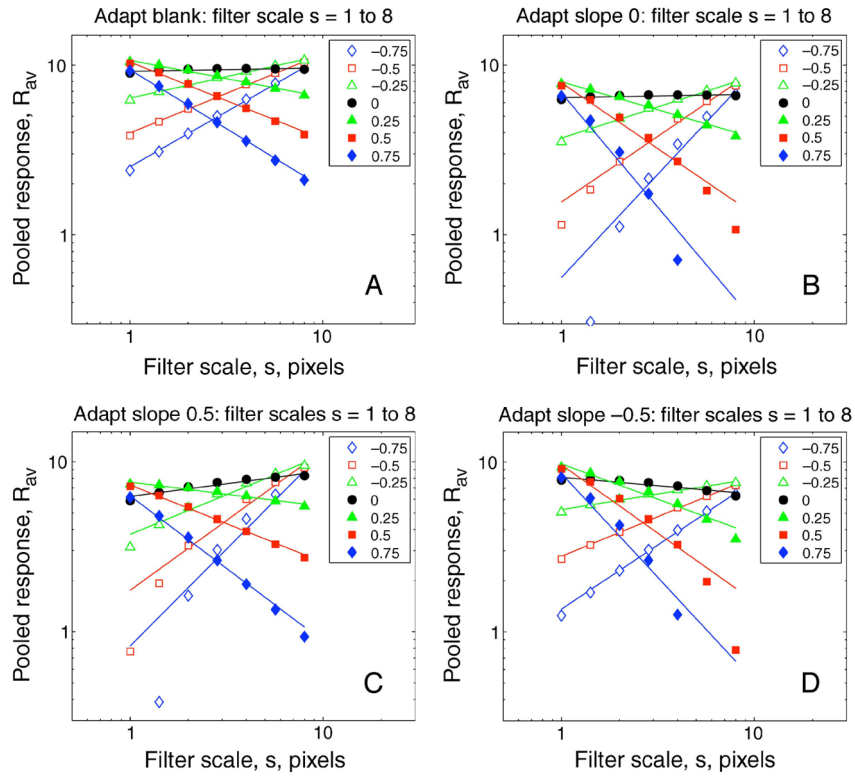
Figure 7.
How the simple (subtractive adaptation; no norm) model's pooled response R_{av} varies with filter scale s after adaptation ($k = 0.3$) to (A) blank, (B) $1/f$ spectrum ($\alpha = 0$), (C) sharpened ($\alpha = 0.5$), or (D) blurred ($\alpha = -0.5$) image spectra. Images were spatial noise, as in Experiments 1 and 2. Symbols represent different test slope offsets (a) as shown, and slope of fitted lines indicates the resulting blur code b. Note that increasing filter scale s corresponds to *decreasing* the preferred spatial frequency of the filter. Compared with (B), the fitted lines tend to rotate anticlockwise, meaning greater blur, after sharpened adaptation (C), but clockwise, meaning greater sharpness, after blurred adaptation (D).
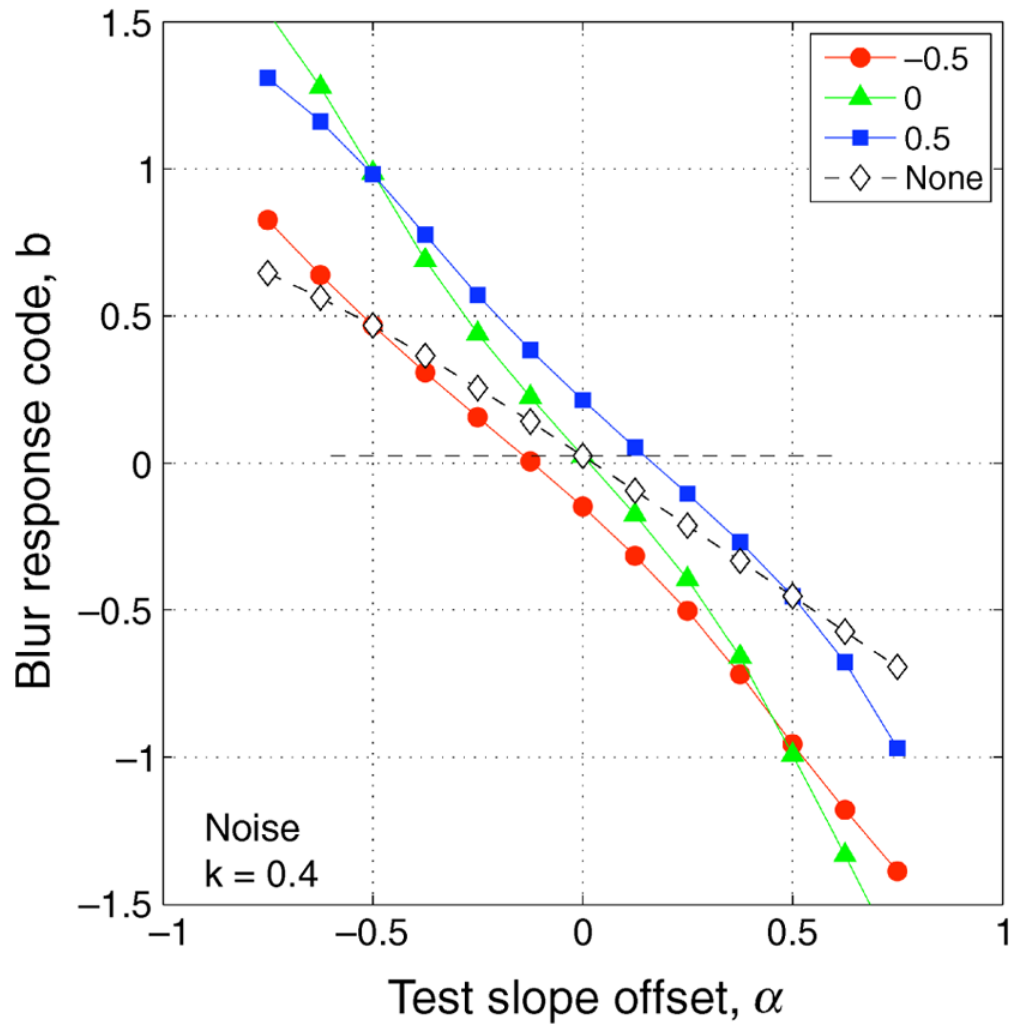
Figure 8.
Blur codes b as a function of test slope offset α, for noise images, after subtractive adaptation ($k = 0.4$) to blurred ($\alpha = -0.5$, red circles), $1/f$ ($\alpha = 0$, green triangles), or sharpened ($\alpha = +0.5$, blue squares) noise images, or no subtraction (open diamonds).

Figure 9.
(A) Observed and predicted blur matches for the noise images in Experiment 1. Predictions are from the simple fatigue model (Equation 2), with adaptation strength $k = 0.3$. (B) Blur aftereffects across the range of test blurs, expressed as the difference between observed and veridical matches, for sharpened (blue squares, $\alpha = 0.5$) and blurred (red circles, $\alpha = -0.5$) adapters. The model exhibits "repulsion" of test blur matches away from the adapting blur (dashed circles) for all 3 adapters. This was a very poor prediction for Experiment 1.
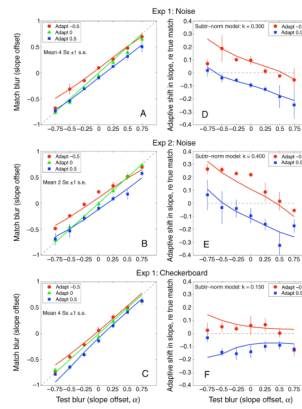
Figure 10.
(A–C) Observed blur matches for different test blurs compared with predictions of the norm-based model for the 2 experiments reported in this paper. Adaptation strength (*k*) was chosen to fit the data separately in (A), (B), and (C). (D–F) Blur aftereffects (from (A)–(C)), expressed as the difference between observed and veridical matches, for sharpened adapters (squares, 0.5) and blurred adapters (circles, −0.5). For Experiment 2, predictions of the fatigue model are the same as the norm model shown here, but for Experiment 1, they are very different (and incorrect), as seen in Figure 9.

Figure 11.
Predictions of the norm-based model ($k = 0.5$) compared with the blur aftereffects of Webster et al. (2002), derived from judgments of best focus after adaptation to natural images with various levels of blur or sharpening. The fatigue (no-norm) model makes identical predictions for this experiment. Oblique dashed line shows where data would lie if full normalization took place, i.e., if blurred or sharpened adapters came to look perfectly focused during the period of adaptation. Dashed curve shows model predictions for $k = 0.35$; see text.

Figure A1.
Predictions of a simple divisive gain control (Equation A1, with $m = 0.10$) applied to data of Experiment 2. See Appendix A for details. This model shows a uniform, but partial, normalization of blur matches across all test blurs. Compared with Figures 10B and 10E, the fit is poor when adapt and test blurs are very different.

Figure A2.
Fits of a more elaborate gain control model (Equation A2, with $k = 0.20$) applied to data of Experiment 2. See Appendix A for details.