An unusual kind of contrast adaptation: Shifting a contrast comparison level

S. Sabina Wolfson

Department of Psychology, Columbia University, New York, NY, USA

Norma Graham

Department of Psychology, Columbia University, New York, NY, USA

We have found an unusual kind of contrast adaptation in human pattern vision that seems fundamentally different from previously reported effects. As the observer adapts to different levels of contrast, the visibility of some contrast-defined (second-order) patterns dramatically increases and that of others dramatically decreases. Oddly, visibility is poor for patterns containing contrasts both above and below the recent average contrast. To explain these effects, we hypothesize a new kind of process acting in concert with a known contrast-gain control of the normalization type. The new process compares current contrast to an adaptable comparison level; this level reflects the recent average contrast. Such a process existing at an early stage of visual processing is likely to have widespread effects at higher stages.

Introduction

There are many dimensions along which the visual system adapts (e.g., luminance and color). Without adaptation, the system would function very poorly. We have recently discovered a kind of adaptation to the contrast of visual patterns that is dramatic in magnitude and has qualitative characteristics very different from those previously reported. (In presenting work at conferences, we sometimes call this kind of contrast adaptation "Buffy adaptation." The origin of that term is described in Graham & Wolfson, 2007).

Sensory and perceptual adaptation processes occur at time-scales that range from many minutes (e.g., dark adaptation, contingent aftereffects like the McCollough effects) to fractions of a second, short enough to happen within a fixation (e.g., Muller, Metha, Krauskopf, & Lennie, 1999). The effect we present here is relatively rapid (1 s adapting duration).

The experiment sketched in Figure 1 demonstrates this new kind of contrast adaptation. The observer adapts for 1 s to a grid of identical Gabor patches all at some contrast, for example, 50%. Then, the observer views a brief test stimulus (for 94 ms). The test stimulus is composed of Gabor patches at two different contrasts in alternating rows (or columns), producing contrast-defined stripes that are either horizontal (as in Figure 1) or vertical. Then, the observer views the same adapt stimulus again for 1 s. The observer identifies the orientation of the contrast-defined stripes in the test stimulus.

Figure 1 shows three types of test stimuli that are particularly telling. In the BELOW test stimulus, the two contrast values that produce the stripes, 25% and 45%, are both below the adapt contrast of 50%. In the STRADDLE test stimulus, the contrasts straddle the adapt contrast, and in the ABOVE test stimulus, they are both above it. In each of these test stimuli, the contrast difference (between the two contrast values that produce the stripes) is always the same (20%). In general, observers perform very poorly on the STRADDLE test stimulus and very well on both the ABOVE and BELOW test stimuli.

This was an unexpected result and implicates an unusual kind of contrast adaptation that adjusts a contrast-comparison level based on the recent average contrast.

Methods

All observers were Columbia University undergraduates with normal (or corrected-to-normal) visual acuity. They were paid for their participation. Observer R.K. is listed twice ("rk1" and "rk2") since she ran the whole experiment twice (the second time, intermixed with sessions from additional conditions not reported here).

The experiments were run on a Macintosh G4 with an Iiyama VisionMaster Pro 451 CRT and an ATI Radeon 8500 Mac edition video card. The resolution was 1,280 × 1,024 pixels at 85 Hz. The mean luminance was about 50 cd/m^2. The monitor's look-up-table was linearized. Stimuli were generated and presented using MathWorks'
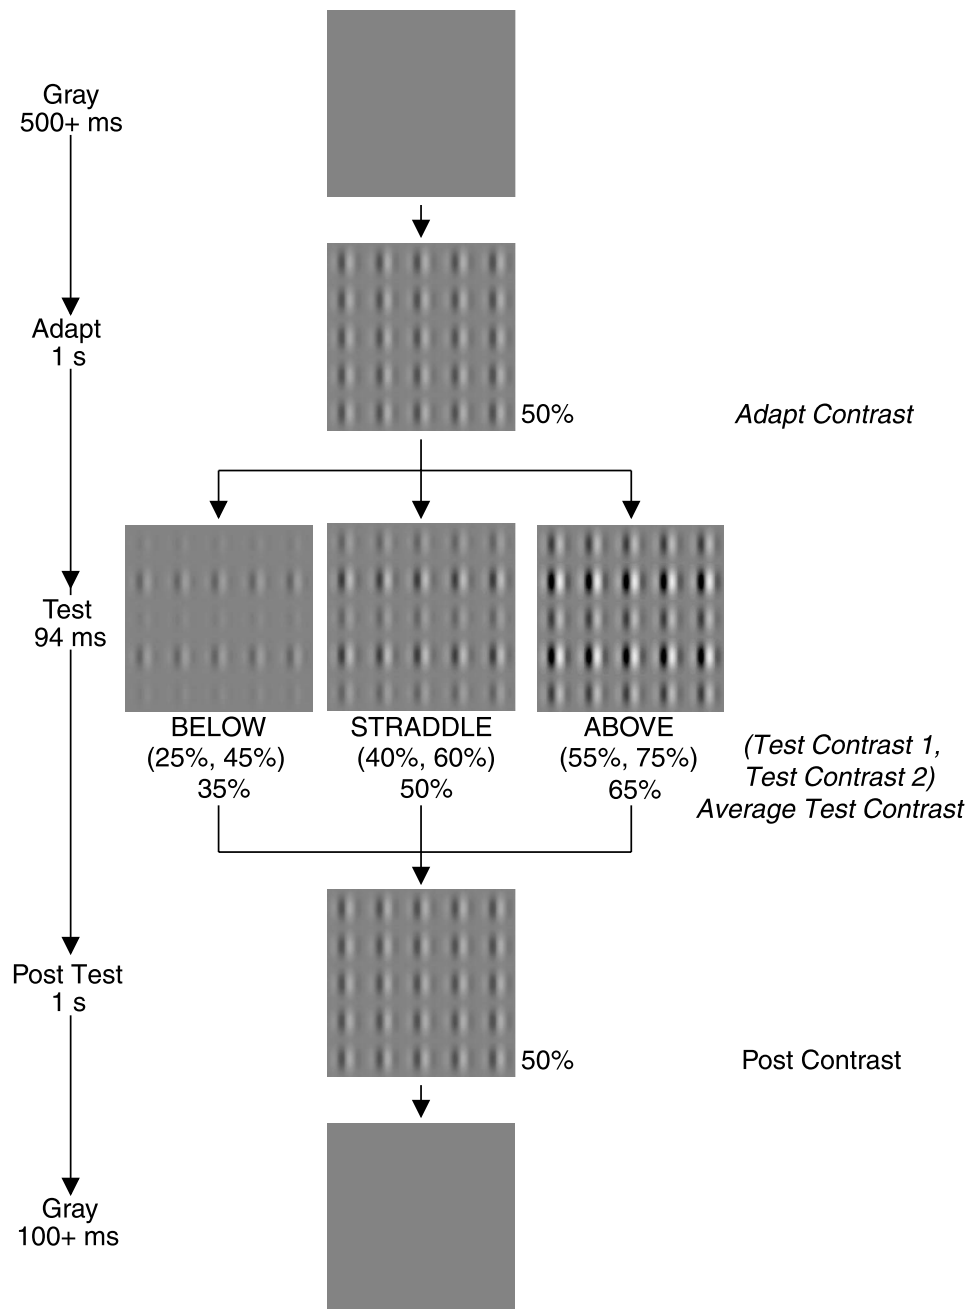
Figure 1. One adapt stimulus and three possible test stimuli are illustrated here as 5 × 5 grids of Gabor patches. (The stimuli used in the experiment were 15 × 15 grids.). The time-course of the experiment is drawn at the left side. Contrast differences in the gray-level illustrations shown here were exaggerated to increase their salience. The contrast of the adapt stimulus (*Adapt Contrast*) is the contrast of the Gabor patches in that stimulus (50% in this example). The contrast values in the test stimulus (*Test Contrast 1*, *Test Contrast 2*) are the contrasts of the two kinds of Gabor patches making up the test stimulus and are enclosed in parentheses on the figure. The *Average Test Contrast* is the average of these two values.

MATLAB with the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997).

Each stimulus was a 15 × 15 grid of Gabor patch elements. (A Gabor patch is a sinusoidal grating windowed by a two-dimensional Gaussian function.) Each Gabor patch was truncated at 1 × 1° (64 × 64 pixels) at the viewing distance of 90 cm. (Distances are approximate as observers'

heads were not constrained.) The center-to-center Gabor patch distance was 64 pixels. The sinusoidal grating in our Gabor patches had a period of 0.5° (32 pixels), which is a spatial frequency of 2 c/deg. In each Gabor patch, a positive zero-crossing of the sinusoidal grating was centered under the Gaussian function. The Gaussian function had a full-width-at-half-height of 0.5° (32 pixels). The contrast of a

Gabor patch is computed by taking the difference between the luminance at the peak of the Gaussian and the mean luminance of the pattern and then dividing that difference by the mean luminance.

Each trial proceeded as follows: The observer pressed the "0" key to start the trial, the screen was gray for 500 ms, the adapt pattern was shown for 1 s, the test pattern was shown for 94 ms, the adapt pattern was shown again for 1 s, the screen was gray for 100 ms, and then the computer beeped, indicating that the observer had to respond. The screen remained gray between trials. The mean luminance was constant throughout the experiment. The observer's task was to identify the orientation (vertical or horizontal) of the contrast-defined stripes in the test pattern using the computer's keyboard. Feedback was provided. The room was dark.

There were four different test pattern configurations: (1) horizontal (second-order) stripes composed of horizontal (first-order) Gabor patch elements, (2) horizontal stripes of vertical elements (as in Figure 1), (3) vertical stripes of vertical elements, and (4) vertical stripes of horizontal elements. The orientation of the Gabor patches was always the same throughout a trial.

There were three adapt contrasts: 35%, 50%, and 65%. The difference between the contrasts of the two element types in the test patterns was always either 10% or 20%. Within a session, trials of all combinations of adapt and test stimuli were intermixed. Not all test stimuli were used with all adapt contrasts in this particular experiment. For example, the data points in Figure 2 show all combinations of test stimuli and adapt contrast used with a 20% contrast difference.

Each session was 320 trials long. Each subject ran at least nine sessions, and thus, each point in Figures 2 and 3 represents at least 72 trials.

Results

Figure 2 shows results (from an experiment like that in Figure 1) after adapting to three different contrast levels (35% in red, 50% in black, and 65% in blue). The vertical axis shows the percentage correct identification of the orientation of the contrast-defined stripes of the test stimulus. The horizontal axis shows the *average* of the two contrasts in the test stimulus. The *difference* between these two contrasts is always 20%. Performance is severely impaired when the average test contrast equals the adapt contrast. For example, each 35% adapt contrast (red) curve has a minimum at an average test contrast of 35%, which corresponds to a STRADDLE test stimulus composed of 25% and 45% contrast Gabor patches. And performance is near perfect to the right of the minimum (ABOVE test stimuli) and to the left of the minimum (BELOW test stimuli). This same pattern of results (poor performance for the STRADDLE test stimulus and very good performance on ABOVE and BELOW test stimuli) is seen for all three of the adapt contrasts.

Another way of thinking about the results in Figure 2 is to consider all the data points directly above a particular average test contrast on the horizontal axis, for example, 65%. The observers were near perfect on this test stimulus (containing 55% and 75% contrast Gabor patches) if they had previously adapted for 1 s to 35% contrast (red) but much worse if they had adapted to 65% contrast (blue). Adaptation so dramatically alters the observer's ability to see contrast-defined patterns that it can change performance on patterns from near perfect to near chance, or vice versa, depending on what contrast the observer has recently seen.
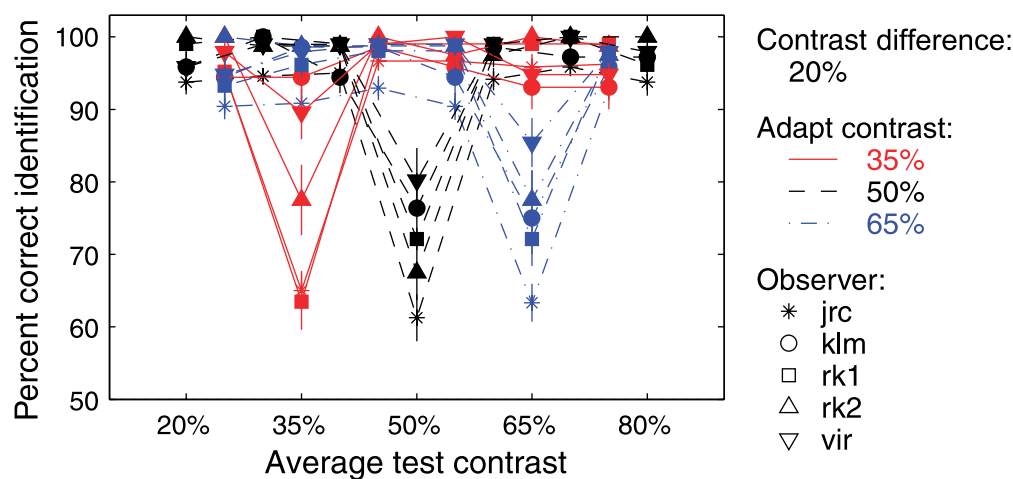

Figure 2. Results from an experiment like that illustrated in Figure 1. The difference between the two contrast values in the test stimulus was always 20%. Observers (indicated by different symbols) adapted to a 15 × 15 grid of identical Gabor patches of either 35% (red solid line), 50% (black dashed line), or 65% (blue dash-dotted line) contrast. Error bars show ± 1 *SEM* across sessions. Performance is very good on ABOVE and BELOW test stimuli and poorer on STRADDLE test stimuli.
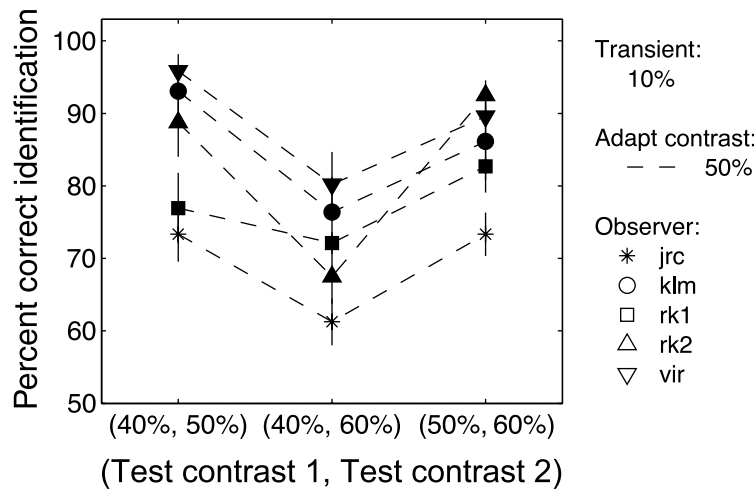
Figure 3. Results from stimuli in which the transient the absolute value of the maximum change between adapt stimulus and test stimulus was always 10%. Other conventions as in Figure 2. Performance is still poor on STRADDLE test stimuli and very good on ABOVE and BELOW test stimuli, even when the transient is held constant.

Some readers may have worried that, while the difference between the two contrasts in the test stimuli in Figure 2 is always 20%, this difference might not be the "same" in all cases in the following sense. Consider the *absolute value of the maximum change* (called the *transient* below) between the two contrasts in a test stimulus and the adapt contrast. In particular, consider a 50% adapt contrast. For the STRADDLE test stimulus, which consists of contrasts 40% and 60%, the transient is just 10%. However, the transient is 20% for the ABOVE test stimulus composed of 50% and 70% contrasts and also 20% for the BELOW stimulus composed of 30% and 50% contrasts. It is greater than 20% for any other ABOVE or BELOW test stimulus.

Figure 3 shows the results for three test stimuli in which the transient (the absolute value of the maximum change) is always 10%. The pair of contrast values in each test stimulus is shown on the horizontal axis. The adapt contrast was 50%. The vertical axis shows percent correct identification.

Clearly, even when the transient is the same, observers perform more poorly on the STRADDLE test stimulus than on the ABOVE and BELOW test stimuli.

Discussion

What kind of visual process could produce results like those in Figures 2 and 3?

It is straightforward to show that rectification on a luminance dimension (as in conventional complex or second-order channels—Graham & Sutter, 1998; Landy & Graham, 2003; Schofield, 2000) cannot produce results like those of Figures 2 and 3, nor can contrast-controlled adaptation processes of the types often called (see, e.g., Ibbotson, 2005) contrast-gain controls (sketched in Figure 4,
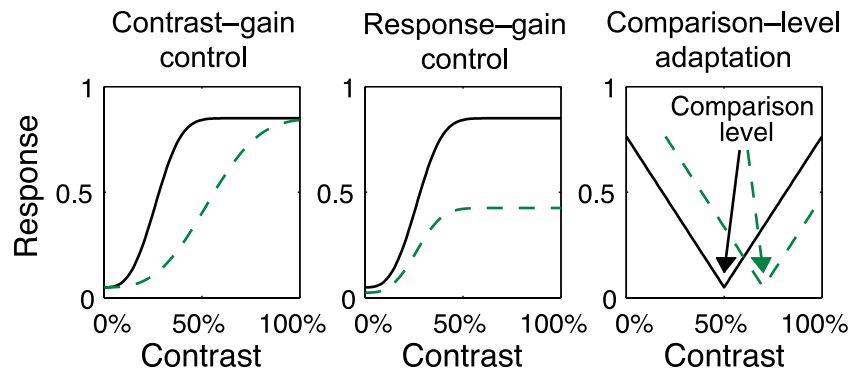

Figure 4. Contrast–response functions from three different kinds of contrast-controlled adaptation mechanisms. The solid line in each panel is the curve after adaptation to one level of contrast. The dashed line is the curve after adaptation to another higher level of contrast. Monotonic contrast–response functions have been assumed to shift generally horizontally (left panel) and/or vertically (middle panel) by adaptation processes. In the adaptation mechanism we propose (right panel), a non-monotonic rectifying function moves horizontally, changing the position of the bottom of the "V" (the comparison level) to equal the recent average contrast.

left panel) or response-gain controls (Figure 4, middle panel) since neither can produce the selectively poor performance in the STRADDLE conditions. (One can consider the functions in Figure 4 as characterizing individual neurons, or as characterizing a suitable group of neurons, or as a more abstract description of a process in a psychophysical model.)

An adaptable comparison process operating on the dimension of contrast (Figure 4, right panel) can explain our results, at least qualitatively. In this process, the function relating response to contrast at each image spatial position is a non-monotonic rectifying function. The contrast value at the minimum of the function will be called the *comparison level*. Adaptation moves the function horizontally by updating the comparison level at a spatial position to equal the recent time-averaged contrast in some neighborhood around that position. Thus, the output of this process is the *un-signed* difference between the current contrast at a position and the comparison level there. Increments in contrast produce approximately the same outputs as decrements of similar magnitude. In this sense, increments and decrements are confusable by the comparison process.

If the function in Figure 4 (right panel) were indeed a perfect full-wave rectifying function, then performance by this mechanism on STRADDLE test stimuli would be at chance, no matter how high the contrast difference, a result we have not found with any observer we tested. Overcoming this and other difficulties with the idea in Figure 4 (right panel) can be accomplished in a number of different ways while incorporating an adaptable comparison process into standard models of pattern vision that include simple (linear) and complex (second-order) channels sensitive to different ranges of spatial frequency and orientation. For example, one can assume that the function shown in Figure 4 (right panel) is not a perfect full-wave rectification but something between a half-wave and a full-wave linear rectification (where the asymmetry can be in different directions in different channels). Or one can assume a function that is not even piecewise linear (due, perhaps, to an early logarithmic transformation). Alternately, one might assume that there are not only channels incorporating this new process but also channels that do not. All these various possibilities have testable consequences. We are currently investigating the quantitative success of such extended models and trying to distinguish among them.

Why has this comparison-level process not been suggested before?

We think that this kind of adaptation effect was not noticed earlier because experimenters, in general, have not tested pattern discriminations like those here after adapting to different contrasts. For example, all of our own past experiments using this kind of pattern discrimination only

adapted to 0% contrast (blank gray field), so we would not have seen this adaptation.

To our knowledge, the only other situations in which visual adaptation introduces "confusion" or "lack of discriminability" between values on either side of the adapting value are situations in which the adaptation occurs not on the contrast dimension but on high-level dimensions like those describing face perception (Rhodes et al., 2005; Rhodes, Maloney, Turner, & Ewing, 2006).

Why has this form of adaptation evolved?

Most of the functions suggested for perceptual adaptation (e.g., in Clifford & Rhodes, 2005) belong in one of two classes, and to some extent, both classes may apply to the adaptation phenomena here:

1. To re-center the operating range of the system to be at or near the current adaptation level (the average level in the recent past of whatever kind of input is at issue) so that performance is optimized near that level. The function of light adaptation is widely believed to be of this sort.
2. To suppress the response to unchanged visual stimuli and thereby highlight the responses to changes because changes signal important events in the environment and/or to make neural coding more efficient.

Consider the first class of explanation: An operating range seems to be moving in Figure 2 as the adapt contrast changes, but the movement seems to make performance worse near the adapting level, not better. However, not illustrated in Figure 2 is one important fact well established from our prior work: Without any previous adaptation to pattern contrast (more exactly, after adaptation to a blank gray field, i.e., to 0% contrast), performance on most of the test stimuli in Figures 2 and 3 would be very poor, but performance would be very good on test stimuli of even lower average test contrast than that plotted. Indeed, in the absence of adaptation to non-zero contrast, performance for the test stimuli plotted at the right end of Figures 2 and 3 would be close to or at chance (e.g., Graham, Beck, & Sutter, 1992; Graham & Sutter, 2000; Wolfson & Graham, 2005). Thus, adaptation to a non-zero pattern contrast of 35%, 50%, or 65% in Figure 2 can be said to move the operating range to the right relative to that without pattern adaptation, thus producing better performance than before on the test patterns near the adapting level (except for STRADDLE stimuli). Results after adaptation to a blank gray field (0% contrast) have been successfully explained by incorporating into the model a contrast-gain control of the normalization type which acts on the outputs of both simple (first-order) and complex (second-order) channels (e.g., Graham et al., 1992; Graham & Sutter, 2000). The adaptable comparison process proposed here and that previously-identified contrast-gain control (of the

normalization type) work together. The consequence of their combined action is that most patterns near the adapt contrast (all except for STRADDLE patterns) are easy to perceive, but patterns composed of contrasts far away from the adapt contrast are difficult to perceive.

One might describe the result here (Figures 2 and 3) in words appropriate to the second class of proposed function: As a consequence of the action of the adaptable comparison process, the visual system is very sensitive to (most) contrast changes from the adapting level to new but nearby contrast levels (the test levels). These changes can be used by the system to identify features like the orientation of contrast modulation here, as long as the changes do not straddle the initial contrast. In short, the system responds to change well, consistent with this second class of proposed function, but loses information about the sign of the change.

Thus, both classes of explanation may provide some understanding of why this kind of contrast adaptation exists, but neither class helps us understand the very poor performance on STRADDLE test stimuli. Perhaps, wiring a neural system so that it can signal a change quickly without regard to sign is much less costly (in terms of whatever kinds of costs that limit evolution of neural tissue) than wiring a system to signal quickly both a change and its sign. If so, we do not understand why it might be so.

Or, perhaps, there is some evolutionary advantage to *not* being able to perform well on those stimuli. But, if so, it is a mystery to us.

Where in the nervous system might this adaptable comparison process exist?

Single neurons in cortical area V1 have been extensively studied (see, e.g., Carandini et al., 2005), but nothing like this adaptable contrast-level comparison process has ever been reported. Gardner et al. (2005) looked at human fMRI BOLD responses from populations of neurons in V1, V2, V3, and hV4 to changes in test stimulus contrast. They found that V1, V2, and V3 responses are positive to increments in test stimulus contrast from adapting contrast and negative to decrements. Responses in hV4, however, are positive to either increments or decrements in contrast from an adapting contrast. This "confusion" of increments and decrements is consistent with a possible role for hV4 in the comparison process of Figure 4 (right panel). Further, human fMRI data from Larsson, Landy, and Heeger (2006) strongly suggest that stimulus orientation in second-order tasks such as ours is extracted largely in VO1 (but also V3A/B and LO1), all areas beyond V1 and V2.

Consequences

The possible consequences of contrast adaptation like that proposed here (Figure 4, right panel) have not been considered previously. It is worthwhile doing so. Any adaptation that occurs at relatively low levels of visual processing will affect the later processes of visual perception. The adaptation described here may be particularly important for the perception of shape and form.

Acknowledgments

Commercial relationships: none.
Corresponding author: S. Sabina Wolfson.
Email: sabina@psych.columbia.edu.
Address: Department of Psychology, Columbia University, 406 Schermerhorn Hall, New York, NY 10027.

References

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 443–446. [PubMed]

Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., et al. (2005). Do we know what the early visual system does? *Journal of Neuroscience, 25,* 10577–10597. [PubMed] [Article]

Clifford, C. W., & Rhodes, G. (2005). *Fitting the mind to the world.* Oxford, UK: Oxford University Press.

Gardner, J. L., Sun, P., Waggoner, R. A., Ueno, K., Tanaka, K., & Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. *Neuron, 47,* 607–620. [PubMed] [Article]

Graham, N., Beck, J., & Sutter, A. (1992). Nonlinear processes in spatial–frequency channel models of perceived texture segregation: Effects of sign and amount of contrast. *Vision Research, 32,* 719–743. [PubMed]

Graham, N., & Sutter, A. (1998). Spatial summation in simple (Fourier) and complex (non-Fourier) texture channels. *Vision Research, 38,* 231–257. [PubMed]

Graham, N., & Sutter, A. (2000). Normalization: Contrast-gain control in simple (Fourier) and complex (non-Fourier) pathways of pattern vision. *Vision Research, 40,* 2737–2761. [PubMed]

Graham, N., & Wolfson, S. S. (2006). Complex channels become more complex: Modeling a contrast adaptation process [Abstract]. *Journal of Vision, 6*(6):694, 694a, http://journalofvision.org/6/6/694/, doi:10.1167/6.6.694.

Graham, N., & Wolfson, S. S. (2007). Exploring contrast-controlled adaptation processes in human vision (with help from Buffy the Vampire Slayer). In L. R. Harris & M. R. M. Jenkin (Eds.), *Computational vision in neural and machine systems* (pp. 9–47). Cambridge, UK: Cambridge University Press.

Ibbotson, M. R. (2005). Physiological mechanisms of adaptation in the visual system. In C. W. G. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world* (pp. 15–46). Oxford, UK: Oxford University Press.

Landy, M. S., & Graham, N. (2003). Visual perception of texture. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1106–1118). Cambridge, MA: MIT Press.

Larsson, J., Landy, M. S., & Heeger, D. J. (2006). Orientation-selective adaptation to first- and second-order patterns in human visual cortex. *Journal of Neurophysiology, 95,* 862–881. [PubMed] [Article]

Muller, J. R., Metha, A. B., Krauskopf, J., & Lennie, P. (1999). Rapid adaptation in visual cortex to the structure of images. *Science, 285,* 1405–1408. [PubMed]

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. [PubMed]

Rhodes, G., Maloney, L. T., Turner, J., & Ewing, L. (2006). Is the average face special [Abstract]? *Journal of Vision, 6*(6):283, 283a, http://journalofvision.org/6/6/283/, doi:10.1167/6.6.283.

Rhodes, G., Robbins, R., Jacquet, E., McKone, E. Jaffery, L., & Clifford, C. W. G. (2005). Adaptation and face perception: How aftereffects implicate norm-based coding of faces. In C. W. G. Clifford & G. Rhodes (Eds.), *Fitting the mind to the world* (pp. 213–240). Oxford, UK: Oxford University Press.

Schofield, A. J. (2000). What does second-order vision see in an image? *Perception, 29,* 1071–1086. [PubMed]

Wolfson, S. S., & Graham, N. (2005). Element-arrangement textures in multiple objective tasks. *Spatial Vision, 18,* 209–226. [PubMed] [Article]