

To appear in C. Chubb, B. Doshier, Z. Lu, and R. Schiffrin (Eds.), *Vision, Memory, and Attention*, American Psychological Association (2012).

Two Visual Contrast Processes: One New, One Old

Norma Graham and S. Sabina Wolfson

Dept. of Psychology, Columbia University, NY, NY 10027

nvg1@columbia.edu, sw354@columbia.edu

In everyday life, we occasionally look at blank, un-textured regions of the world around us, a blue unclouded sky for example. But most of the time our eyes see regions occupied by spatial patterning – by texture, form or patterns -- as when looking at a person to whom we are talking or at the text on this page. Further, there is constant temporal change as well as spatial patterning – if only as a result of eye movements. Thus, the eye is usually looking at a visual scene where different parts of the scene are characterized by different levels of visual contrast, and, from moment to moment, the contrast at any point on the retina is changing. (Visual contrast in any region of the scene is the difference between the lightest and darkest parts of that region, relative to some measure of overall intensity in that region.) So one might wonder how the spatial patterning an observer has just seen affects the visual processing of the spatial patterning that an observer sees now. And, more specifically, one might wonder how the visual contrast one has just seen in a region affects the processing of visual contrast there now.

The first part of this chapter is about an effect of contrast adaptation discovered rather recently, nicknamed Buffy adaptation. (For the origin of the nickname, see Graham and Wolfson, 2007. We are using the term adaptation here only to mean the effect of preceding contrast on the processing of subsequent visual contrast. Our procedure, which will be described in Fig. 1, might also be called masking or a procedure to study temporal processing.) This recently discovered effect of contrast adaptation dramatically increases the visibility of some contrast-defined patterns and dramatically decreases that of others. The second part of the chapter briefly places this new effect in the context of a previously known effect (called the old effect here), which exhibits more conventional Weber-law-like behavior.

THE NEW EFFECT AND HYPOTHESIZED PROCESS

Place Figure 1 about here

Figure 1 shows a typical trial from the experiments reported here. The observer looks at a gray screen briefly, and then sees a pattern we will call the Adapt pattern. The duration of the Adapt pattern is approximately one second except where noted. Then, for an even briefer time, approximately 100 ms except where noted, the observer sees a Test pattern, which can be one of a number of different possibilities depending on the experiment. Three possibilities are illustrated in Fig. 1. The Test pattern is followed by a Post-Test pattern that here is exactly the same as the Adapt pattern. The screen then returns to gray. (The mean luminance of the screen stays constant throughout the experiment.)

The Adapt pattern is a grid of Gabor patches all of the same contrast. The images in Figure 1 show 5x5 grids of Gabor patches, but the experiments reported here used either 15x15 or 2x2 grids. The Test pattern on any trial here is exactly like the Adapt pattern except that the contrasts of the Gabor patches will generally differ between Adapt and Test. Each Test pattern is composed of two potentially different contrasts of Gabor patches (which vary from trial to trial); the Gabor patches are arranged so that the two contrast values define a striped pattern.

For the experiments reported here, the spatial frequency of the sinusoidal fluctuation in the Gabor patches is 2 cycles per degree; there is approximately one full cycle in each patch. In the grid there is one row of patches per degree of visual angle, also one column per degree of visual angle. The Gabor patches are either all vertical or all horizontal, changing randomly from trial to trial. The contrast values shown in the figure are illustrative only, and those for each study will be specified separately. (Each contrast value is the contrast of a Gabor patch, expressed as one-half the difference between peak and trough luminances divided by the mean luminance underneath the Gaussian window of the Gabor patch.) The contrast of the Adapt pattern remains constant throughout a block of trials, but the test contrasts vary from trial to trial.

The observer's response is to indicate whether the contrast-defined stripes in the Test pattern are "horizontal" or "vertical", as illustrated in the right half of Figure 1. Feedback is provided as to the correctness of the response.

After responding the observer initiates the next trial with a key press so the duration of gray on the screen between one trial's Post-Test pattern and the next trial's Adapt pattern is significantly longer than 600ms required by the trial's characteristics. The gray screens were at the same mean luminance as the patterns.

The new result: the straddle effect

Figures 2 and 3 show results of a study with trials like those in Figure 1 with the following contrast choices: the contrast of the Adapt pattern could be 35%, 50%, or 65%; the two contrasts in the Test pattern always differ by 10% for the results shown here but their average varies. For further details of the study see the figure legends and Wolfson & Graham (2007b).

Place Figure 2 about here

Figure 2 shows the results from several observers for a subset of 3 kinds of trials. In all three, the Adapt pattern's contrast is 50%, and the difference between the two contrasts in the Test pattern is 10%. But the two test contrasts vary (values are shown in the parentheses) and can be above the adapt contrast (top row), or can straddle it (middle row), or can be below it (bottom row). The performance was much worse for the Straddle test pattern (53%, 57%, and 61% correct) than for the Above or Below test pattern (between 86% and 99% correct).

According to one common point of view, the function of adaptation (e.g. light adaptation) along a dimension (e.g. light intensity) is to re-center the operating range of the system to be at or near the current adaptation level (the recent time-averaged level on the dimension). This re-centering is done so that the system can respond near optimally in the current environment, responding best for values near the current adaptation level and worse for values farther away. On this point of view, there is something very surprising about the results in Figure 2. The observers perform worst on the Straddle test pattern, which is the Test pattern with contrast values nearest the adapt contrast level. They do better on the Above and Below test patterns, and those are patterns which contain contrast values further away from the adapt contrast level. This is exactly opposite to what is expected from the common point of view mentioned above. And the effect is very large – a point that will be developed further in the next two figures.

Place Figure 3 about here

Figure 3 shows the results for all three adapting contrasts, each paired with many Test patterns. The difference between the two test contrasts was always 10%. As is clear from the V-shaped curves, for any adapt contrast, performance is lowest when the average test contrast equals the adapt contrast (that is, when the test contrasts straddle the adapt contrast). And performance is much better when the average test contrast is substantially above or substantially below the adapt contrast (within the range shown).

Place Figure 4 about here.

In another study (Graham & Wolfson, 2007) we measured contrast-difference thresholds rather than just measuring performance for a fixed test contrast difference. Some results from this study are re-plotted in a new manner in Figure 4. Each panel shows the results for one observer with one duration of test pattern (82 ms in the left column and 35 ms in the right column). The adapt contrast was always 50%. Each point shows the results from trials where the average test contrast was constant at some value in the range from 37.5% to 62.5%, and the difference between the two contrasts varied. The threshold was that contrast difference leading to criterion performance correct for the identification of the orientation of contrast-defined stripes, the task shown on the right of Figure 1. As expected from the percent correct performances in the study shown in Figures 2 and 3, the thresholds measured in this study for the Straddle test pattern (center point on each curve) are substantially higher than those for the Above or Below test patterns. This finding can be represented numerically by the ratio of the Straddle threshold to the other thresholds' average, and this number is indicated on each panel as "th. ratio." The highest threshold ratio was 5.7, and even the lowest ratio was above 2.0. (The black lines fitted to the data points in each panel, and the values of k and g , are from a model that will be discussed later.)

Explanations that do not work for the straddle effect

Consider three patterns, which, like those in Figures 2 and 3, have a difference between test contrasts of 10%. In the Straddle test pattern (composed of 45% and 55% contrasts) the magnitude (absolute value) of the difference between the adapt contrast and either test contrast is only 5%. In the Above pattern composed of (50%, 60%) and the Below pattern composed of (40%, 50%), the magnitude of one of the differences between one test contrast (50%) and the adapt contrast (50%) is zero but the other is 10% (that is, $|60\%-50\%|$ or $|40\%-50\%|$), much larger than the 5% differences in the Straddle pattern. A number of people have asked us whether these larger change magnitudes (lets refer to them as transients) that occur in the Above and Below cases relative to the Straddle case account for the better performance on the Above and Below test patterns. The answer to that question – by a straightforward empirical test – turns out to be “No.” Again using an adapt contrast of 50% consider the following set of three test patterns constructed so that this change magnitude is always 10%: the contrast values in the Below test pattern are (40%, 50%), in the Straddle pattern are (40%, 60%), and in the Above pattern are (50%, 60%). As it turns out empirically, performance in the Straddle case is still substantially worse than in the Above and Below cases. We have now collected results from many such sets of three patterns, or constant-transient trios, and they consistently show much worse performance on the Straddle case. See Figure 3 of Wolfson & Graham (2007) and Figure 7 of Wolfson & Graham (2009).

A nonlinear transducer (monotonic function applied locally) also cannot account for the new effect. This is true even if the transducer is shaped to produce so-called *pedestal effects* for

both increments and decrements in contrast. The discussion is rather lengthy so we do not include it here. The interested reader can see *Appendix A1: Shifting monotonic transducer* in Wolfson & Graham (2009).

An explanation of the straddle effect: Shifting, rectifying, contrast comparison

How can one make sense of the experimental results? It is as if, within whatever system is responsible for the observer's performance in these experiments:

- (1) Adaptation to contrast resets some comparison level to represent the recent time-average contrast seen at each place in the visual field;
- (2) And the current contrast at each spatial position pattern is evaluated relative to that comparison level,
- (3) with increases and decreases in contrast being equally salient but quite confusable with each other (within whatever system is doing this perceptual task).

Place Figure 5 about here

This idea is diagrammed as the comparison process illustrated in Figure 5. The core of this idea is that there exists something in the visual system that acts like a rectification function on the contrast dimension. This rectification makes it hard to perceive a Straddle test pattern correctly because, when the Straddle test pattern appears after the Adapt pattern, the increase in contrast produces the same output as the decrease. Thus, it is difficult for observers to identify the orientation of the contrast-defined stripes in the Straddle test pattern. Indeed, if the rectification function were a full-wave function, like that in Figure 5, then increases and decreases of the same magnitude would produce identical outputs from this comparison process and thus be totally confusable by the subsequent stages of visual processing. The full-wave rectification function shown is too extreme to quantitatively fit our results since it would predict that an observer's performance would be at chance on all Straddle test patterns, no matter how large the difference between the two contrasts. We have never seen an observer quite that extreme. This aspect of Figure 5 is modified below.

Conventional second-order (complex, *FNF*, non-Fourier) channels are composed of two linear-filter stages with a rectification-type nonlinearity between the two stages. These channels can be described informally as structures in which neurons with small receptive fields feed into neurons with large receptive fields with a point-wise rectification in-between. Thus these channels are selectively sensitive to both the first-order spatial-frequency and orientation content (determined primarily by individual Gabor patches) and the second-order spatial-frequency and orientation (determined by overall arrangement of contrast in the grid of Gabor patches). These second-order channels are a very useful addition to the previously-suggested simple (first-order, Fourier) channels and can account for many aspects of human pattern perception as has been shown by many different investigators in a large number of studies. (See Graham 2011 for a recent overview of this work with many references.)

In Graham & Wolfson (2007), Figures 2.22 and 2.23 with accompanying text present a channel model that is modified from the conventional second-order channel model. The modified channel contains, in addition to the components of the conventional channel, a process with behavior like that in Figure 5. To explain the experimental results for any observer we have yet studied, we cannot use a strict full-wave rectification for that process as mentioned above. We need some compromise between a half-wave and a full-wave rectification in this channel model. We embody this compromise in our current model by using pairs of channels. The

rectification functions in any pair of channels are mirror-symmetric, each somewhere between a full-wave and a half-wave rectification, so that one channel is more sensitive to contrast increases and the other more sensitive to contrast decreases. The degree of compromise between a half- and a full-wave rectification functions can be represented by a parameter k , which varies from 0 for half-wave to 1 for full-wave. More completely, its value equals the absolute value of the ratio of the slope of the rectification function's shallower side to the slope of its steeper side.

An alternate approach would be to consider mixtures of channels some having strict full-wave and some having strict half-wave rectification functions (e.g., Sperling 1989, Sperling Chubb Solomon & Lu 1994). We have not tried this yet, but we suspect the predictions from a mixture model would look very similar to those presented here and that the parameter k values here would be monotonic with the ratios of full-wave to half-wave channels in the best-fitting mixture model's predictions.

Lets look at the threshold results in Figure 4 again. The black lines fitted to the data points show best-fitting predictions from a model including adaptable complex channels. The simple model we used to generate these predictions and produce the fits here assumes that all stages of visual processing were deterministic (there was no noise in the responses). It also assumes that the channel with the maximum output determined the observer's response. Finally, it is a static model (there is no explicit representation of time) but the assumption is made that adapt pattern duration is long enough to make the comparison level equal to the adapt contrast. See Graham 2011 for the equations. What is represented by the rectification function on the contrast dimension in this simple static model might well be unpacked in a dynamic model to be something that came from the temporal processing of contrast transients (or perhaps even of luminance transients although that seems unlikely to be true in the visual system).

The model predictions (black lines) fitted to the data points in Figure 4 result from a 2-parameter fit: k is the ratio of slopes referred to before, and g is an overall gain parameter. The fit is excellent.

There are individual differences in both parameters. The sensitivity parameter g varies substantially among observers, which is not very surprising. It is determined in the fits by the thresholds to the Above and Below test patterns (the thresholds on the horizontal line segments at the left and right ends of the predicted curves). The parameter k is more interesting. It varies substantially among datasets. It was determined in the fits by the threshold ratio "*th.ratio*" (the ratio of straddle thresholds to the above and below thresholds). A value of $k = 0$ (half-wave rectification) corresponds to a threshold ratio of 2. Something close to this value occurs for one dataset. A value of $k = 1$ (strict full-wave) corresponds to a threshold ratio of infinity (the threshold for the Straddle would be infinitely high) and did not occur in our results. The highest values of this parameter k were near 0.6 and indicate rectification functions middling between half- and full-wave.

THE OLD EFFECT AND HYPOTHESIZED PROCESS

We switch now to a situation that does *not* show the new straddle effect that we have been discussing so far. Instead this situation shows an old effect that is an example of generalized Weber-law behavior. All the results shown so far in this chapter involved middle-range adapt contrasts and middle-range average test contrasts. On the basis of our own previous work (e.g., Graham & Sutter 2000, Wolfson & Graham 2005) and of extrapolations from others'

work, we already knew a great deal about the effects of one other, very extreme, adapt contrast – namely the effect of an adapt contrast of 0%.

The old result: Weber-Law behavior

Adapting to 0% contrast means adapting to a steady homogeneous gray field with no Gabor patches in it. Such a condition is often called “no adaptation” or “before adaptation” and may be the condition most often used in experiments on pattern perception. We have studied this condition a great deal ourselves in the past, and have referred to it this way. We no longer view it this way for reasons we hope will be clear by the end of this chapter.

Place Figure 6 about here

Let's look at some results of adapting to 0% contrast. For this study (done recently to replicate the older results but with the current observers) we used test patterns in which the test contrast difference was 10%, and the average test contrast took on closely spaced values from 10% to 90%. This is almost, but not quite, the full range possible. The smallest and largest possible average test contrasts for patterns characterized by a 10% test contrast difference are 5% and 95% respectively. (There is no Test pattern that can straddle an Adapt pattern of zero percent contrast or one of 100% contrast.) As shown in the top part of Fig. 6, performance after adapting to 0% contrast is very good for the lowest average test contrasts used – that is, for test contrasts near the adapt contrast. Performance then drops monotonically (within the limits of experimental variability) as the average test contrasts gets far away from the adapt contrast. This decrease in performance is reminiscent of many effects seen in perception. It and its close relatives have been much studied. We will refer to this behavior as Weber-law-like behavior, or just Weber behavior for short, for the reason indicated in Figure 6 bottom part. The leftmost of the three example Test patterns there has an average test contrast of 10% and contains two contrasts of 15% and 5%, producing a contrast ratio of 3. The Test pattern in the middle is composed of 55% and 45%, and the contrast ratio is 1.2. The Test pattern on the right has a still smaller contrast ratio. Thus, in general, performance gets worse as the ratio of the higher contrast to the lower contrast in the Test pattern gets smaller. In previous more extensive experiments, we have further shown that – after adaptation to 0% contrast – it is indeed this ratio that predicts performance over a large part of the range. For example, if you considered another Test pattern with the same ratio, i.e. 60% and 20% compared to 15% and 5%, then the observer's performance would likely be very similar on both (e.g. Graham & Sutter, 1996, 2000; Wolfson & Graham, 2005). Dependence on the ratio of intensities occurs on many dimensions (here the dimension is contrast) and is a generalization of the behavior described in Weber's law.

As it turns out, the Weber behavior in this situation *cannot* be explained by a nonlinear monotonic function applied locally (e.g. Fechner). For brevity's sake we do not describe this here. (An interested reader could consult Graham & Sutter, 2000 or Wolfson & Graham 2009.)

An explanation of the Weber behavior: inhibition among channels in a normalization network

A model that can account for the Weber-law-like behavior after adaptation to 0% contrast includes inhibitory interconnections among simple (linear) and complex (second-order) spatial-frequency and orientation-selective channels. The inhibitory interactions are in a form often called normalization, a form of divisive suppression in which the response of each neuron is divided by (normalized by, has its contrast gain set by) the total output of a pool of neurons. This kind of model provides an excellent fit for a wide variety of visual patterns. (Recent discussions of this kind of normalization process can be found in Reynolds & Heeger, 2009 and Graham, 2011.) In particular, this kind of model correctly predicts the Weber-like behavior after

adaptation to 0% contrast for a large range of patterns like those used here (e.g. Graham & Sutter 2000).

THE OLD AND THE NEW TOGETHER

Place Figure 7 about here

What happens for those combinations of adapt and test contrasts that we have not yet discussed? In particular, what happens when you adapt to middling contrasts and test with very low or very high average contrasts? And what happens when you adapt to very high contrasts? Figure 7 shows the results of an experiment (Wolfson & Graham, 2009) with five different adapt contrasts ranging from 0% (in the bottom panels) to 100% (in the top panels) and pairs of test contrasts that varied from the lowest possible to the highest possible (their average is plotted on the horizontal axis). The difference between the test contrasts was always 10%.

Adaptation to 0% (bottom panels) replicates earlier results (e.g. Figure 6). There are individual differences in sensitivity; SYP (on left) is a very sensitive observer (highest performing for a given contrast difference) while RK (on right) is less sensitive.

What happens after adaptation to a middle contrast for the full range of average test contrasts (e.g. 50% in the middle row of Figure 7)? We again see the bad performance on the Straddle test patterns with good performance for Above and Below test patterns of contrasts somewhat outside the Straddle range (as in Figure 3). But now we also see bad performance for Test patterns at the very ends of these curves, having contrasts very far above or very far below the adapt contrast. To put it another way, the observers' best performance overall is for test contrasts at an intermediate distance from the adapt contrast.

As you look from the top panel of Figure 7 (adapt contrast of 100%) down to the bottom panel (adapt contrast of 0%), you can see movement of the region of peak performance from right to left, following the adapt contrast, where this region of peak performance has a notch at the adapt contrast itself. (This is clearest for the three intermediate curves where it is possible to have an average test contrast equal to the adapt contrast, but the notch is also clear at an adapt contrast of 100%.) So the range of contrast patterns that are perceived best – at least in the sense of the perception required in our experimental task – generally shifts to regions near the adapting contrast. This shift is consistent with a frequently suggested function of adaptation: that of moving the operating range to suit the current conditions. The exception to the general shift is the very bad performance on the Straddle test patterns, which are patterns with average test contrasts right in the middle of the good-performance region.

We suspect that these curves in Figure 7 result from the interaction of two processes across the full range of contrasts but dominating in different ranges. The adaptable contrast-comparison-level process (nicknamed Buffy adaptation) is the primary cause of the bad performance for the Straddle. The normalization process produces the Weber-like behavior – where the dimension is the unsigned difference between the test contrast and the adapt contrast. And thus it occurs both for increments and decrements in contrast (both for average test contrasts substantially above and for those substantially below the adapt contrast).

We have recently started some fitting of these results with predictions from a model that combines the normalization model we used previously for adapting to 0% (instantiated in very simple equations) and the contrast-comparison process (also instantiated in very simple equations). The results are encouraging (Graham, 2011).

We have been interested in studying the dynamics of both the straddle effect and of the Weber behavior to help uncover the properties of the presumed processes (contrast comparison and normalization), their neural substrate, and also their functionality in human vision. We have been doing two kinds of study. The first kind used flickering adapt patterns of various temporal frequencies (Graham and Wolfson 2007). The second kind used different durations of adapt patterns in the situation of Figure 1 here. Results for one observer in one version of this second kind of study are shown in Figure 8.

Place Figure 8 about here

The results of both kinds of dynamic experiment suggest the following conclusions which should be held tentatively given the absence to date of any modeling of these dynamics: The re-setting of the comparison level in the new process may happen largely if not completely within 50 ms to 100 ms; The temporal integration of the contrast-gain control in the old normalization process takes a second or longer.

DISCUSSION

The duration of a typical eye fixation is approximately 200-250 ms, and the duration of a typical saccade is approximately 50-70 ms. Thus the hypothesized resetting of the contrast-comparison level might well occur within an eye fixation or even within a saccade. Such fast adaptation may be important since (quoting Frazor and Geisler, 2006) "...eye movements are frequently large enough that there will be little correlation in the contrast or luminance on a receptive field from one fixation to the next, and thus rapid contrast and luminance gain control are essential."

Is the perceiving of changes (perceiving the change in contrast between the Adapt and Test patterns in the experiments here) so desirable that loss of direction-of-change information (loss which leads to bad performance on Straddle test patterns in the experiments here) can persist as a side effect throughout the course of evolution? And, unless it is positively advantageous to perform badly in Straddle situations, which seems unlikely, the question arises: Could the visual system perfectly track the sign of the change as well as the fact of the change? If the answer is "yes" why doesn't it here? If the answer is "No", why is the answer "No"? We are beginning to suspect that there is some design reason why it is best (fastest, easiest, cheapest) to detect changes in a way that partially loses information about sign (direction of change). Perhaps the design constraint has something to do with some difficulty neural networks have in responding stably to rapid changes (transients) since large imbalances from different parts of the network are likely to occur at transients. Other possible design constraints are suggested in recent studies that ask whether some features of adaptation (in particular, orientation-selective adaptation) may occur as a reaction to changes in an organism's internal state rather than to changes in the external environment. They consider factors like sparseness and variance of cortical spike trains, and they explore these factors' implications for optimal processing at different time scale. Perhaps some such design constraint could help explain the puzzling loss of direction-of-change information in the Straddle effect here.

References

- Frazor, R. A., & Geisler, W. S. (2006). Local luminance and contrast in natural images. *Vision Research*, *46*, 1585-1598.
- Graham, N. (in press). Beyond multiple pattern analyzers modeled as linear filters (as classical V1 simple cells): Useful additions of the last 25 years. *Vision Research*. doi10.1016/j.visres.2011.02.007
- 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.
- Graham, N. & Wolfson, S. S. (2007). Exploring contrast-controlled adaptation processes in human vision (with help from Buffy the Vampire Slayer). In Laurence R.M. Harris & Michael R. Jenkin (Eds.) *Computational Vision in Neural and Machine Systems* (pp 9-47). United Kingdom: University Press, Cambridge.
- Reynolds J. H. & Heeger D. J. (2009). The Normalization Model of Attention. *Neuron*, *61*, 168-185.
- Sperling, G. (1989). Three stages and two systems of visual processing. *Spatial Vision*, *4*, 183-207.
- Sperling, G., Chubb, C., Solomon, J. A. & Lu, Z.-L. (1994). Full-wave and half-wave rectification in motion and texture. In *Higher-order processing in the visual system Ciba Symposium No. 184* (pp. 287-308). Chichester: Wiley.
- Wolfson, S. S. & Graham, N. (2005). Element-arrangement textures in multiple objective tasks.

Acknowledgements

This work was supported in part by National Eye Institute Grant EY08459. We thank Boris Grinshpun, Ian Kwok, Stephanie Pan, Alisa Surkis, and Jiatao Wang, all of whom did modeling which is not discussed explicitly here but which has been very important for the development of the ideas. We also thank our observers for their hours of effort.

And we would like to express years of gratitude to George Sperling for many contributions, and particularly for his consistent use of rigorous mathematical models that are complicated enough to capture interesting behavior but simple enough to give understanding and insight into the processes producing the behavior.

Figure Legends

Figure 1. A typical trial from the experiments reported here. The experimental set-up is described somewhat further in Wolfson & Graham (2007b). The conditions shown in this figure will be referred to as the standard conditions, and exceptions will be explicitly noted when they occur. Note that only one Test pattern is shown on any trial, but 3 example Test patterns are shown on this figure.

The right half illustrates the task of the observer. On each trial the observer indicates whether the contrast-defined stripes are horizontal or vertical. Two examples of patterns corresponding to each of these responses are shown.

Figure 2. Results from one Adapt pattern followed by each of several Test patterns – a subset of the results in Figure 3. Performance is given as percent correct identification of the orientation of the contrast-defined stripes where each percentage is from approximately 50 trials of a given trial type. (15x15 grids. Spatial characteristics and timing are standard as in Figure 1 and accompanying test. 3 observers: AG, CG, and CT.)

Figure 3. The results of adapting to different contrasts 35%, 50%, and 65% (identified by arrows on the upper and lower horizontal axes) are shown by three different line types. The three observers are represented by different symbol types. The difference between the two contrasts in the Test pattern is always 10%. From the same study as Figure 2.

Figure 4. Contrast-difference thresholds as a function of average test contrast. The adapt contrast was 50% contrast. The mean \pm 2 standard errors across sessions is plotted at each point. This full set of thresholds is unpublished but done under the same conditions as those in Graham & Wolfson (2007). See text for description of the threshold ratio *th. ratio* and the parameters *k* and *g* used to produce the model fits to the data points (black lines). The observer was KF for both panels in the top row. The other three observers at the longer test duration (left column) were JW, EG, AG from second panel down. The other two observers at the shorter duration (right column) were KN and SH from the second panel down.

Figure 5. Diagram of a process that rectifies with a shifting “zero-point” – the comparison level – that adapts to equal the average of its recent input (the average recent local contrast). The output of this possible comparison process (at any particular position) is proportional to the un-signed difference between the current input contrast and the comparison level. Thus, under suitable temporal conditions, the output is proportional to the magnitude of the change from the Adapt pattern to the Test pattern, but it ignores the direction of the change.

Figure 6. The top shows performance after adapting to 0% contrast for Test patterns in which the contrast difference was 10%. Results are shown for average test contrasts from 10% to 90%. Patterns were 2x2 grids of Gabor patches. 3 observers are indicated by different symbols: MM (stars), RK (squares), and VR (upside-down triangles).

The bottom shows the ratio of the higher test contrast to the lower test contrast in three example Test patterns. The performance of the observer is clearly correlated with this ratio.

Figure 7. Results of experiment covering the full range of adapt and test contrasts. The five different adapt contrasts are marked in the right side and indicated by diamond symbols on the horizontal axis. They range from 0% (in the bottom panels) to 100% (in the top panels). The difference between the test contrasts was always 10% and the average varied from 5% to 95% (the lowest-possible to the highest-possible). The patterns were 2x2 grids. These data are from 2 of the 3 observers shown in Wolfson & Graham (2009).

Figure 8. The effect of varying the Adapt pattern duration between 0 and 1000 ms (shown by the right-hand labels) when adapting to 50% contrast (indicated by double vertical line). The duration of the test pattern was 82ms. The post-Test pattern duration equaled that of the Adapt pattern. The duration of the gray screens (the first and last events in the trial as shown in Figure 1) were such that the total length of the trial was the same for all adapt durations. There difference between the two contrasts in a Test pattern was always 10%. The patterns were 2x2 grids. Observer SYP.