7

SPATIAL CHARACTERISTICS OF A CONTRAST-COMPARISON PROCESS

S. Sabina Wolfson and Norma Graham
Columbia University

We are interested in the "hidden" intermediate visual processes that are difficult to study because they are far from the retina (light stimulus) and also far from the observer's behavior. There are good reasons to think that there are many intermediate processes and that they are important for perception. These processes were an important part of Naomi Weisstein's interests (see, for example, Weisstein & Harris, 1980; Weisstein et al., 1975).

In the long ages of evolution, eyes occasionally saw (and still occasionally see) a blank, unchanging part of the visual field, for example, a large region of blue sky. But most of the time we are looking at regions of the visual field occupied by spatial texture or forms or patterns. This spatial patterning is almost always changing in time, if only as the result of eye movements. So one might well wonder how the spatial pattern that one has just seen affects the visual processing of the spatial pattern that one sees now. In the course of trying to study something else, we found an effect of this kind – an effect of the preceding pattern's contrast on the current pattern's perceived appearance – that surprised us. This effect of contrast adaptation dramatically increases the visibility of some contrast-defined patterns and dramatically decreases that of others. Initially we half-jokingly referred to this as "Buffy adaptation" as a placeholder because we did not yet understand it. (For the origin of the nickname, see Graham & Wolfson, 2007.) This chapter continues our study of this effect of contrast adaptation.

Figure 7.1 shows a typical trial from the experiments reported here. The observer looks at a fixation screen for a second, then a gray screen for about 500 msec, and then a pattern we will call the *adapt pattern*. (This could also be called a mask pattern as Foley, 2011, does. We briefly discuss the overlapping relationship of the terms "adaptation" and "masking" in the Introduction section of Wolfson & Graham, 2009.) The duration of the adapt pattern is 1 sec. The adapt pattern is followed immediately by a *test pattern*. The duration of the test pattern is approximately 100 msec. The test pattern is followed immediately by a *post-test pattern* that is exactly the same as the adapt pattern. The screen then returns to gray (for at least 100 msec) until the observer makes a response.

The adapt pattern is a 2x2 grid of Gabor patches all of the same contrast. (The contrast of a Gabor patch is expressed as one-half the difference between peak and trough luminances divided by the mean luminance underneath the Gaussian window of the Gabor patch.)

Each test pattern is composed of two different contrasts of Gabor patches (which vary from trial to trial). The observers' responses indicate whether they think the contrast-defined stripes in the test patterns are "horizontal" or "vertical", as illustrated on the right side of Figure 7.1. Feedback is provided as to the correctness of the response on each trial. After an observer responds, the screen remains gray until the observer initiates the next trial with a key press.

The spatial frequency of the sinusoidal fluctuation in the Gabor patches is always 2 cycles/degree in the work we report here. There is approximately one full cycle in each Gabor patch. The mean luminance of the screen stays constant throughout the experiment at about 50 cd/m^2.

In the *baseline condition* in our experiments, the test pattern on any trial is exactly like the adapt pattern (and post-test pattern) except for Gabor-patch contrasts. This is the condition illustrated in Figure 7.1.

Many details of the experiments and patterns are given in Table 7.1. The table lists experimental parameters that were not necessarily the same for each experiment. The following parameters were the same for all experiments in this chapter: grid size = 2x2 Gabor patches; Gabor patch spatial frequency = 2 cycles/degree; viewing distance = 90 cm with unrestrained head; adapt contrast = 50 percent; adapt pattern's duration = 1 sec; test pattern's duration = 94 msec; post-test

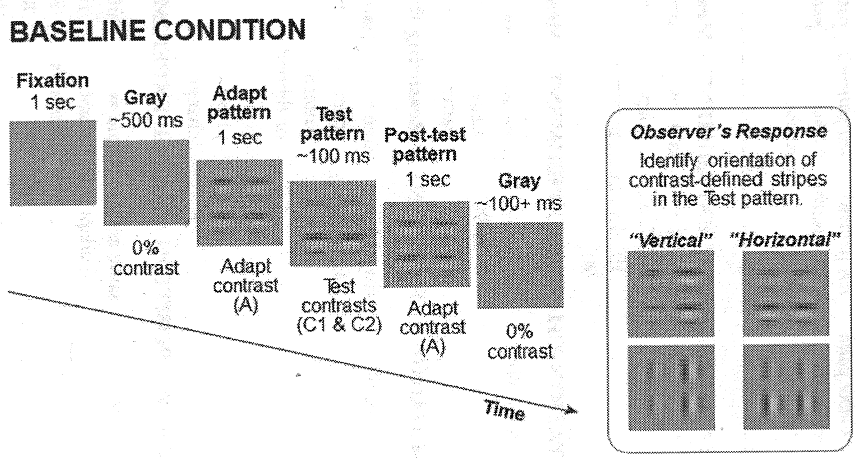

FIGURE 7.1 A typical trial in the *baseline condition*. See text for description.

TABLE 7.1

Plotted in…	*Adapt pattern's Gabor patch orientation (Øadapt)*	*Orientation of test pattern's Gabor patches*	*Phase of test pattern's Gabor patches*	*Position of test pattern's Gabor patches*	*Test contrast differences*	*Average test contrasts |C1-C2|*	*Center-to-center Gabor patch distance*	*Observers*	*Trials per session*	*In house expt name*
ORIENTATION-CHANGE EXPERIMENT - Measuring performance										
Figure 4	0° or 90°	Øadapt +/− 0°, 2°, 5°, 10°, 15°, or 90°	same as adapt pattern's	same as adapt pattern's	10% 20% 30%	one Above, one Straddle, and one Below for each |C1-C2|	2 deg	LG MC	216	BG1orient4i89
ORIENTATION-CHANGE EXPERIMENT - Measuring threshold										
Figure 5	0° or 90°	Øadapt +/− 0° or 90°	same as adapt pattern's	same as adapt pattern's	varied to determine threshold	40% 50% 60%	1 deg	LG BSG LG BSG	90	HGP4i101b H4i101b
PHASE-CHANGE EXPERIMENT - Measuring threshold										
Figure 6	0° or 90°	same as Øadapt	same as adapt pattern's or flipped (180° phase change)	same as adapt pattern's	varied to determine threshold	40% 50% 60%	1 deg	MC BSG MC BSG	90	HG1phase4i101b H4i101b
HORIZONTAL-POSITION CHANGE EXPERIMENT - Measuring performance										
Figure 7	0° or 90°	same as Øadapt	same as adapt pattern's	shift to left or right by 0/32, 2/32, 5/32, 8/32, 16/32, 19/32, 24/32, 29/32, or 32/32 cycles	10% 20% 30%	one Above, one Straddle, and one Below for each |C1-C2|	1 deg	LG MC	306	BJ4i85

pattern's duration = 1 sec. Within each experiment, trials from the baseline condition (in which the adapt, test, and post-test patterns are identical except for contrast) and the other conditions were randomly intermixed. In the horizontal-position-change experiment we used all of the Δposition values listed but we did not plot all of them in Figure 7.7. The rightmost column lists the experiment names we used in the lab in order to allow cross-publication comparison when the same data gets used in different analyses.

In the *baseline condition* in our experiments, the test pattern on any trial is exactly like the adapt pattern (and post-test pattern) except for Gabor-patch contrasts. This is the condition illustrated in Figure 7.1.

Three example test patterns are shown in Figure 7.2. The two test contrasts (C1, C2) always differ by 10 percent, but the average test contrast varies. The test pattern in the center has an average test contrast of 50 percent. In this test pattern, the two test contrasts were 45 percent and 55 percent. These test contrasts *Straddle* the adapt contrast of 50 percent. The test pattern on the left has an average test contrast of 45 percent; the two test contrasts are 40 percent and 50 percent. This average test contrast is *Below* the adapt contrast. The test pattern on the right has an average test contrast of 55 percent; the two test contrasts are 50 percent and 60 percent. This average test contrast is *Above* the adapt contrast.

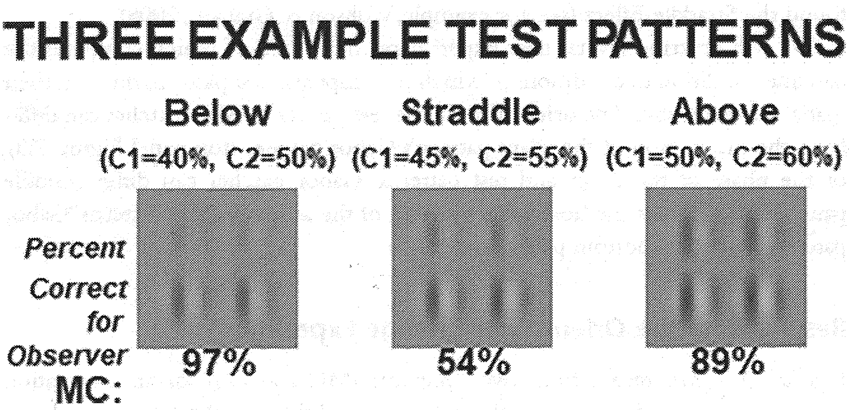

FIGURE 7.2 Three example test patterns, each with a 10 percent test-contrast difference (difference between test contrast C1 and test contrast C2). Grayscale differences have been exaggerated for visual clarity.

The average test contrast of the Below test pattern (45 percent) is "below" the adapt contrast of 50 percent. The average test contrast of the Above test pattern (55 percent) is "above" the adapt contrast of 50 percent. In the Straddle test pattern, the two test contrasts (45 percent and 55 percent) "straddle" the adapt contrast of 50 percent. Performance of observer MC on these test patterns is shown below the grayscale images. While all of these test patterns have the same test-contrast difference, performance on the Straddle test pattern is much worse than performance on the Below or Above test patterns. (The performances here are shown as open circle symbols in the upper left panel of Figure 7.4.)

Result from one observer (MC) in a baseline condition are shown below the test patterns in Figure 7.2 as percent correct. Performance on the Straddle test pattern was worse than for the Below and Above test patterns. We call this effect the *Straddle Effect* and have suggested a contrast-comparison process described later in this chapter to explain this effect.

Foley (2011), using stimuli quite like ours (composed of Gabor patches), also found the Straddle Effect and expanded on our understanding of it. Kachinksy, Smith, and Pokorny (2003), using quite different stimuli (composed of solid squares), and with a different goal (investigating magnocellular vs parvocellular pathways), included contrast conditions similar to ours; if their Figure 6 data is plotted like ours, it shows the Straddle Effect.

Purpose of this Chapter

The purpose of the present chapter is to explore the spatial characteristics of the Straddle Effect in order to elaborate on its hypothesized explanation, the contrast-comparison process.

In all of the experiments reported here, the adapt contrast is always 50 percent and only three kinds of test patterns are used: Below, Straddle, and Above (analogous to those in Figure 7.2). Previously, we have studied the baseline condition with other adapt contrasts and fuller ranges of test contrasts and we have always found the Straddle Effect (see, for example, Wolfson & Graham, 2009).

We will describe several new experiments here, each of which compares the baseline condition to conditions in which the adapt and test patterns differ in their spatial characteristics. The orientation of the test pattern's Gabor patches can differ from the orientation of the adapt pattern's Gabor patches (top panel Figure 7.3); or the phase of the adapt and test patterns' Gabor patches can differ (middle panel Figure 7.3); or the horizontal position of the adapt and test patterns' Gabor patches can differ (bottom panel Figure 7.3).

Results from the Orientation-Change Experiments

Figure 7.4 shows results from two observers (MC and LG) for an orientation change (Δorient) of 0 degrees (the baseline condition) in the leftmost column through an orientation change of 90 degrees in the rightmost column. (An orientation change of 90 degrees is illustrated in the top panel of Figure 7.3.) The open circles (gray squares, black diamonds) show results when the test-contrast difference was 10 percent (20 percent, 30 percent). In each panel, the adapt contrast of 50 percent is marked by an asterisk on the horizontal axis. The points plotted above the asterisk are results from Straddle test patterns. Points to the left are from Below test patterns; points to the right are from Above test patterns.

For each curve in Figure 7.4, the Straddle Effect is clear in the baseline condition: performance on the Straddle test pattern is lower than performance on the

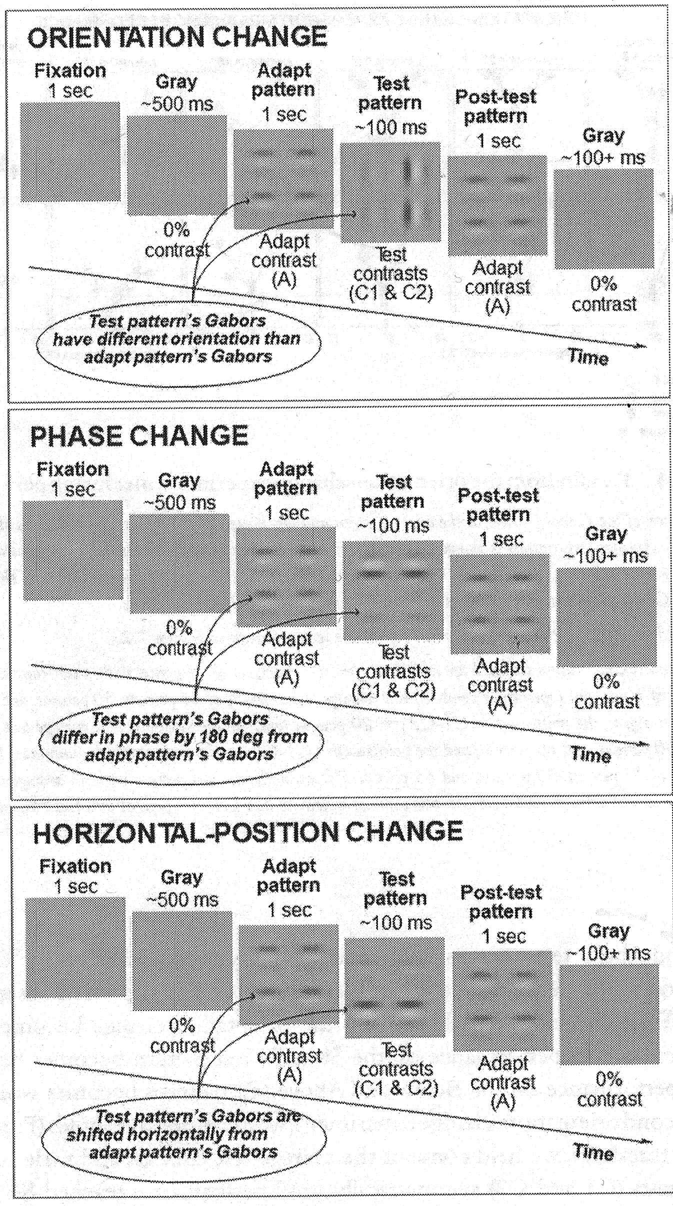

FIGURE 7.3 A typical trial for each of three kinds of experiments: Orientation change (top panel), phase change (middle panel), and horizontal-position change (bottom panel).

These are like the baseline condition except that the test pattern's Gabor patches differ in spatial characteristics from the adapt pattern's Gabor patches. The observer's response – as in the baseline condition – indicates whether he/she thinks the contrast-defined stripes in the test pattern are "horizontal" or "vertical". Feedback (as to the correctness of the response) was provided.

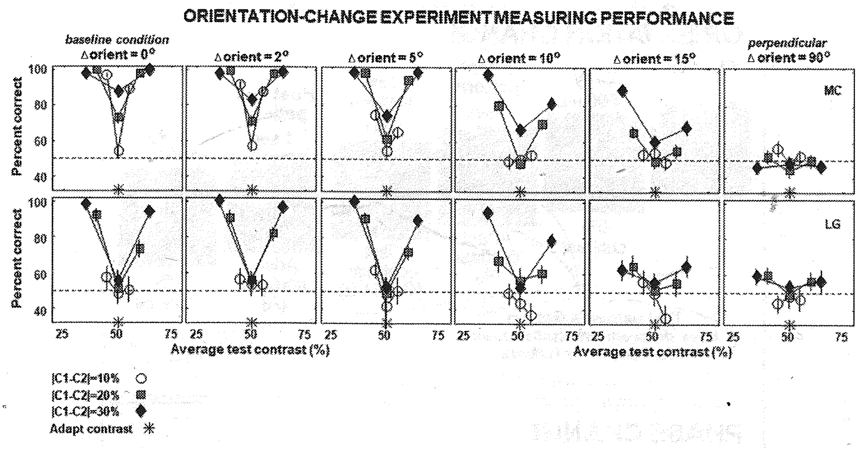

FIGURE 7.4 Results from the orientation-change experiment measuring performance.

The orientation of the Gabor patches in the adapt pattern and test pattern differ by Δorient (which is shown above each column). Average test contrast is shown on the horizontal axis. The asterisk marks the adapt contrast (50 percent) on the horizontal axis. Performance, measured as percent correct, is shown on the vertical axis. Two observers (MC and LG) are shown in the two rows. Error bars show +/− 1 SE across sessions.

The open circle symbols in the upper left panel repeat the results shown in Figure 7.2.

The precise test contrast values we used are as follows: in each panel, the three points with a test-contrast difference (|C1-C2|) of 10 percent (open circle symbols) had average test contrasts of 45 percent, 50 percent, and 55 percent (reading left to right); the points with |C1-C2| = 20 percent (gray square symbols) had average test contrasts of 40 percent, 50 percent, and 60 percent; and the points with |C1-C2| = 30 percent (black diamonds) had average test contrasts of 35 percent, 50 percent, and 65 percent. For example, the test pattern with an average test contrast of 45 percent on the curve for a 10 percent test-contrast difference had C1=40 percent and C2=50 percent.

Below and Above test patterns. (The open circle symbols in the upper left panel correspond to the performance values listed in Figure 7.2.) For each observer, the Straddle Effect becomes less dramatic as the orientation change becomes larger. This is *not* because performance on the Straddle test pattern becomes better but because performance on the Below and Above test patterns becomes worse.

In a second orientation-change experiment, we measured thresholds (Figure 7.5). For each threshold, we held constant the average test contrast and varied the two test contrasts (C1 and C2) symmetrically until performance reached 80 percent correct. The average test contrast was held constant at 40 percent, 50 percent, or 60 percent for the Below, Straddle, or Above test patterns, respectively. Figure 7.5 shows thresholds for the baseline condition (0 degree orientation change, left column) and for the perpendicular condition (90 degree orientation change, right column). Again, the Straddle Effect is present in the baseline condition: the threshold for the Straddle test pattern is higher than those for the Below and Above test patterns. In the perpendicular condition, the Straddle Effect has

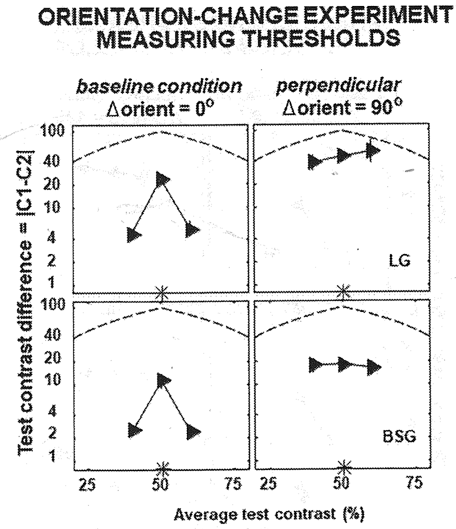

FIGURE 7.5 Results from the orientation-change experiment measuring threshold (rather than performance which is shown in Figure 7.4).

The left column shows results from the baseline condition (orientation change of 0 degrees) and the right column shows results from an orientation change of 90 degrees. Two observers (LG and BSG) are shown in the two rows. The vertical axis shows threshold as the test-contrast difference (|C1-C2|) that leads to 80 percent correct. The horizontal axis shows average test contrast; the three average tests contrasts used were 40 percent, 50 percent, and 60 percent. The asterisk marks the adapt contrast (50 percent) on the horizontal axis. The data points and error bars show the mean +/– 2 SE across sessions. The dashed line shows the maximum possible test-contrast difference.

disappeared because the thresholds on the Below and Above test patterns have become higher.

A Note to Prevent Confusion

Foley (2011) reported an orientation experiment (his Figure 10) which looks superficially like ours, but is in fact a very different experiment. Importantly, he did not keep average test contrast constant while using Below, Straddle, and Above test patterns. Also, his adapt contrast (16 percent) was much lower than ours (50 percent).

Results from the phase-change experiment

In the phase-change experiment, we compared the baseline condition with one phase-changed condition. This phase change – which we call "flipped phase" – is illustrated in Figure 7.3, middle panel; the black and white bars of the Gabor patch switch position between the adapt pattern and test pattern. Figure 7.6, formatted

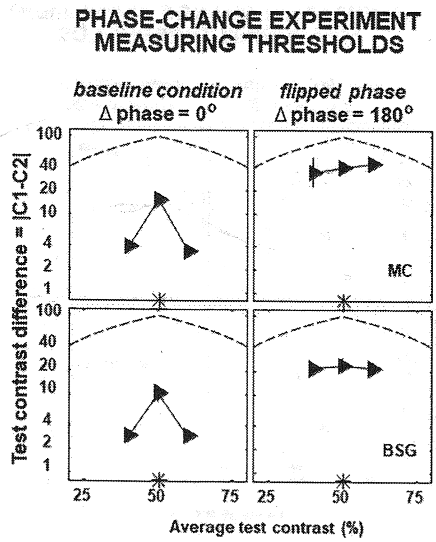

FIGURE 7.6 Results from the phase-change experiment measuring threshold.

The left column is the baseline condition and the right column is the flipped phase condition (in which the phase has been changed by 180 degrees). This figure is in the same format as Figure 7.5.

like Figure 7.5, shows results for the baseline condition (i.e., a phase change of 0 degrees, left column) and for the flipped-phase condition (a phase change of 180 degrees, right column). Again we see the Straddle effect in the baseline condition, but we do not see the Straddle effect in the flipped-phase condition. Again this is because thresholds for the Below and Above test patterns increase to approximately equal that for the Straddle pattern.

Results from the Horizontal-Position-Change Experiment

A horizontal-position change is shown in Figure 7.3, bottom panel; the Gabor patches in the illustrated test pattern are moved left relative to those in the adapt pattern. Figure 7.7, formatted like Figure 7.4, shows results in the baseline condition (leftmost column) and for many different horizontal-position changes (Δposition) up to a maximum of 1 cycle (rightmost column). For each observer, the Straddle Effect becomes less dramatic as the horizontal-position change becomes larger. Again, this is *not* because performance on the Straddle test pattern becomes better but because performance on the Below and Above test patterns becomes worse.

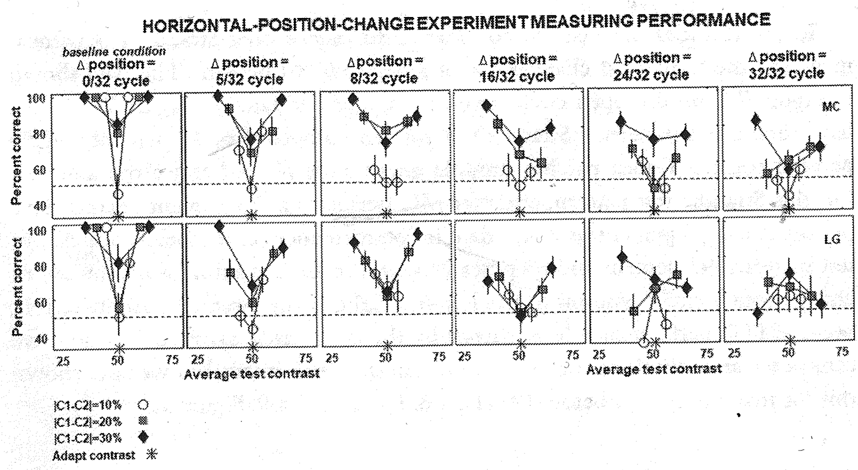

FIGURE 7.7 Results from the horizontal-position-change experiment measuring performance.

The horizontal position of the Gabor patches in the adapt pattern and test pattern differed by Δposition (which is shown above each column). Δposition is measured as fractions of a cycle in the Gabor patches' sinusoidal fluctuation. This figure is in the same format as Figure 7.4.

Note About One Complexity in the Horizontal-Position-Change Experiment

In the baseline condition, the Gabor patches were vertical on half of the trials and horizontal on the other half of the trials. (This was also true in the orientation-change experiments and the phase-change experiment but was not a confounding factor in those experiments.) It might be a confounding factor in the horizontal-position-change experiment because the position change was always horizontal, never vertical. However, as it turns out, results for the vertical Gabor patch trials considered alone look very similar to results for the horizontal Gabor patch trials considered alone. Thus, we have not separated the trials in Figure 7.

Explanation for the Straddle Effect that *Does Not* Work

Consider again the test patterns show in Figure 7.2. The test-contrast difference is 10 percent for each test pattern. However, the *change* in a single Gabor patch's contrast at the transition between the adapt pattern and test pattern is +5 percent or –5 percent for the Straddle test pattern, while the change is 0 percent or –10 percent for the Below test pattern (0 percent or +10 percent for the Above test pattern). Thus one might think that the poorer performance on the Straddle test pattern is because the individual changes are of smaller magnitude (5 percent) than the maximum change on the Below and Above test patterns (10 percent).

To test this idea, we need to consider performance on a Straddle test pattern in which the individual changes are of magnitude 10 percent. The data shown in Figure 7.2 are the open circles from the upper left panel in Figure 7.4, so we have such a test pattern: a Straddle test pattern composed of 40 percent and 60 percent contrast Gabor patches (middle gray square plotted above the asterisk). On this Straddle test pattern, observer MC performs at 73 percent correct. This performance (73 percent) is worse than his performance on the Below and Above test patterns (97 percent and 89 percent) with the same maximum magnitude of change (open circle symbols to the left and right of the asterisk). Unfortunately, because LG's performance level shown by the open circle symbols is so low, we cannot see an example of this in her results in this experiment, but we have shown this for many observers before (Wolfson & Graham, 2009, Figure 7).

Another Explanation for the Straddle Effect that *Does Not* Work

One of the reviewers suggested an alternate explanation that is totally reasonable given the results shown in this paper. Here, only three test patterns (Below, Straddle, and Above) are used. Further, in all these Below and Above test patterns only 2 of the 4 Gabor patches have a different contrast than the adapt contrast; the other 2 Gabor patches have the same contrast as the adapt contrast. The reviewer therefore suggests that perhaps it is the *number* of changing contrast levels, rather than their *signs*, that causes the Straddle Effect. The reviewer suggests that a suitable control would be having the contrast of all 4 Gabor patches in the Above test pattern be above the adapt contrast (similarly, all 4 in the Below test pattern be below the adapt contrast). We have frequently collected such data (e.g. Figure 12 and Figure A4 of Wolfson & Graham, 2009) and these data still show the Straddle Effect (as long as the test contrasts are not all very far below or all very far above the adapt contrast).

Explanation for the Straddle Effect that *Does* Work

The Straddle Effect can be explained by a process (which we call contrast comparison) in which the sign of a contrast change is lost (or very degraded), but the magnitude of that contrast change is preserved. It is as if

- Adaptation to contrast resets some comparison level to equal the recent time-averaged contrast seen at that location in the visual field.
- The current contrast at each spatial position is evaluated relative to its comparison level, with increases and decreases being equally effective but increases being quite confusable with decreases (within whatever visual subsystem is responsible for the observer's performance in our experiments).

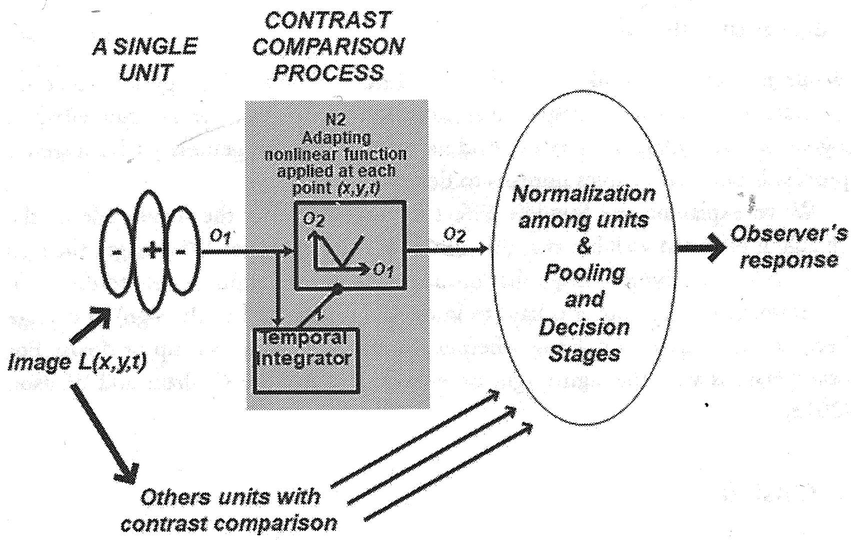

FIGURE 7.8 Our overall model framework showing the shifting, rectifying contrast-comparison process and the normalization process.

The normalization process will not be discussed here; it is a known, important, "hidden" intermediate process (see, e.g., Reynolds & Heeger, 2009, and Solomon & Kohn, 2014). The Pooling and Decision Stages need to include a process or processes that allow the observer to tell the difference between horizontal and vertical arrangements of contrast. This figure adds a typical V1 receptive field as input to the contrast-comparison process within the framework we have shown before (e.g. Figure 17 in Graham, 2011). The addition of a typical V1 receptive field is consistent with the results of this chapter.

Such a contrast-comparison process is diagramed in the gray box in Figure 7.8.

In other words, there exists something in the visual system that acts like a rectification function on the contrast dimension making it difficult to perceive a Straddle test pattern correctly. Consider the extreme case of a full-wave rectification as shown in Figure 7.8. In that case, when a Straddle test pattern appears, the Gabor patches which increase in contrast produce the same output as the Gabor patches which decrease in contrast. The observer *can* see that each Gabor patch has changed, but the observer *cannot* identify the sign of the change in each Gabor patch. Thus, it looks like all four Gabor patches change but all the changes look the same. So the observer cannot see whether the second-order stripes are horizontal or vertical. A full-wave rectification is more extreme than we have found when modeling the results of any observer's performance, but less extreme rectifications have always worked (see Graham & Wolfson, 2013).

For a systematic and careful study of the subjective phenomena in experiments like ours, see Foley (2011).

Helpful or Hurtful?

While it is easy to think of the Straddle Effect as oddly poor performance on Straddle test patterns, we suspect it is more accurately described as a side effect of a generally advantageous process. And we suspect that this generally advantageous process is one that enables humans to detect changes well.

We've explained the Straddle Effect by assuming that the magnitude of the change is encoded well but that the sign of the change is not. Why might the sign be encoded poorly? A change that produces a ringing output is easy to detect in any number of ways, but it is hard to know the direction (i.e., the sign) of change because that requires knowing whether the first excursion was up or down. For other reasons why the sign might be encoded poorly, see Graham and Wolfson (2013).

In Closing

We found that when the spatial characteristics of the adapt and test patterns were identical, there was always a Straddle Effect. But when the adapt and test patterns differed in orientation, phase, or position, the Straddle Effect was much diminished. This diminishing was *not* because performance on the Straddle test patterns improved, but because performance on the Below and Above test patterns got worse.

We currently think that the Straddle Effect is likely to be a side effect of the evolutionary importance of detecting changes/transients.

We initially ran across the Straddle Effect by accident while exploring the dynamics of contrast adaptation using the method that we had previously used to explore the dynamics of light adaptation (Graham & Wolfson, 2007). In those initial experiments we used a large, 15x15 grid of Gabor patch elements, thinking that the Straddle Effect was a global second-order process. As we explored the effect more, we began to wonder if such a large grid was necessary. We tried a 2x2 grid and still found the Straddle Effect, and thus we knew it was more local than we had initially thought. The results reported here suggest that the process producing the Straddle Effect could be as local as a typical V1 receptive field shown in Figure 7.8.

Acknowledgements

This work was supported in part by National Eye Institute Grant EY08459. We thank Ian Kwok and Megan Felder, both of whom did some 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.

Norma would like to express her gratitude to Naomi not only for her published work but for hours of conversation, and Sabina wishes she had had an opportunity to know her.

References

Foley, J. M. (2011). Forward-backward masking of contrast patterns: The role of transients. *Journal of Vision, 11*(9), 15.

Graham, N. (2011). Beyond multiple pattern analyzers modeled as linear filters (as classical V1 simple cells): Useful additions of the last 25 years. *Vision Research, 51*, 1397–1430.

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

Graham, N., & Wolfson, S. S. (2013). Two visual contrast processes: One new, one old. In C. Chubb, B. A. Dosher, Z.-L. Lu, & R. M. Shiffrin (Eds). *Human Information Processing: Vision, Memory, and Attention.* Washington, DC: American Psychological Association.

Kachinsky, E. S., Smith, V. C., & Pokorny, J. (2003). Discrimination and identification of luminance contrast stimuli. *Journal of Vision, 3*(10), 599–609.

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

Solomon, S. G., & Kohn, A. (2014). Moving sensory adaptation beyond suppressive effects in single neurons. *Current Biology, 24*, R1012–R1022.

Weisstein, N., & Harris, C. S. (1980). Masking and the unmasking of distributed representations in the visual system. In C. S. Harris (Ed.) *Visual Coding and Adaptability* (pp. 317–364). Hillsdale, NJ: Lawrence Erlbaum Associates.

Weisstein, N., Ozog, G., & Szoc, R. (1975). A comparison and elaboration of two models of metacontrast. *Psychological Review, 82*(5), 325–343.

Wolfson, S. S., & Graham, N. (2007). An unusual kind of contrast adaptation: Shifting a contrast comparison level. *Journal of Vision, 7*(8), 1–7.

Wolfson, S. S., & Graham, N. (2009). Two contrast adaptation processes: Contrast normalization and shifting, rectifying contrast comparison. *Journal of Vision, 9*(4), 1–23.