PERGAMON

Vision Research 40 (2000) 2737–2761

Vision Research

www.elsevier.com/locate/visres

Normalization: contrast-gain control in simple (Fourier) and complex (non-Fourier) pathways of pattern vision

Norma Graham [a,*], Anne Sutter [b]

[a] *Department of Psychology, Columbia University, New York, NY 10027, USA*
[b] *Department of Psychology, Loyola University, 6525 North Sheridan Rd, Chicago, IL 60626, USA*

Received 13 December 1998; received in revised form 2 November 1999

Abstract

Results from two types of texture-segregation experiments considered jointly demonstrate that the heavily-compressive intensive nonlinearity acting in static pattern vision is not a relatively early, local gain control like light adaptation in the retina or LGN. Nor can it be a late, within-channel contrast-gain control. All the results suggest that it is inhibition among channels as in a normalization network. The normalization pool affects the complex-channel (second-order, non-Fourier) pathway in the same manner in which it affects the simple-channel (first-order, Fourier) pathway, but it is not yet known whether complex channels' outputs are part of the normalization pool.

Keywords: Texture; Inhibition; Contrast gain control; Normalization; Non-Fourier; Light adaptation

1. Introduction

In explanations of perceived texture segregation and related perceptual phenomena, nonlinearities that occur at rather low levels of the visual system have proven to be very powerful. At least two different kinds of nonlinearities — one *intensive* in character and one more intrinsically *spatial* — have been useful (e.g. Sperling, 1989; Malik & Perona 1990; Graham, Beck & Sutter, 1992; Wilson 1993). The primary aim of the present paper is to explore further the heavily-compressive *intensive nonlinearity* that has been implicated in a number of phenomena, particularly texture segregation. We will show that this intensive nonlinearity is achieved through a global contrast-gain control like that in a normalization network. However, understanding this intensive nonlinearity will require considering its interaction with a *spatial nonlinearity*, so we start here by briefly introducing this spatial nonlinearity.

1.1. Spatial nonlinearity: complex channels

Complex channels have been proposed as the mechanism for the spatial nonlinearity and are shown as parts of Figs. 1 and 2. They or similar mechanisms have also been called 'non-Fourier mechanisms' or 'second-order units' or by more specialized terms like 'collector units' or 'collator units'. Complex channels consist of two layers of filtering — where the first is sensitive to higher spatial frequencies than the second — with an intermediate stage consisting of a pointwise nonlinear function of the rectification type. As a consequence of this filter–rectification–filter structure, complex channels respond to low-spatial-frequency patterns of high-spatial-frequency elements. (A simple channel consists of a single filtering stage. For further description of why simple and complex channels respond as they do, see, e.g. Sutter, Beck & Graham, 1989; Graham, Beck & Sutter 1992. For a brief review of the various phenomena for which complex channels or similar mechanisms have been proposed, see introduction to Graham & Sutter, 1998).

One previous finding about complex channels in texture segregation is critical to our arguments below. The results of experiments measuring the tradeoff between the area and contrast of individual elements in element-

* Corresponding author. Tel.: +1-212-6786805; fax: +1-212-8543609.
E-mail address: nvg@psych.columbia.edu (N. Graham).

0042-6989/00/$ - see front matter
PII: S0042-6989(00)00123-1

arrangement textures have implications for the pointwise function at the complex channels' intermediate stage. The results suggest that this function is not piecewise linear (as in conventional full-wave or half-wave rectification) but (in addition to its rectifying action) it is quite expansive. By 'expansive' we mean that its output is a positively accelerating function of its input. It is well described as a power function of the absolute value of the input, where the power is 3 or 4 (Graham & Sutter, 1998).

1.2. Intensive nonlinearity: two candidates

A number of phenomena in perceived texture segregation and related perceptual tasks suggest the existence of another nonlinear process — one that is highly compressive (its output is a decelerating function of its input) and has as its input something closely related to either light intensity or contrast. Two quite different candidates have been suggested for this highly compressive process, which we will refer to in general as the intensive nonlinearity.

1.2.1. Normalization (inter-channel inhibition, global contrast-gain control)

The first candidate (Fig. 1) is inhibition among the channels or some other form of contrast-gain control that depends on a global signal (by which we mean a signal reflecting a rather wide band of spatial frequencies and orientations, but not necessarily a wide area of space). The physiological substrate for this might be intracortical inhibition. In our work we have modeled such inhibitory interaction by a normalization or global-gain-control network, influenced by the work of Bonds (1989, 1993) and Robson (1988a,b) on V1 neurons and specifically by the mathematical models of Heeger (1991, 1992a,b).

In a normalization network, the response of each neuron is divided by (is normalized by, has its contrast gain set by) the total output of a pool of neurons. This approach may not adequately represent all varieties of intracortical inhibition and lateral interaction in V1 (Carandini, Heeger & Movshon, 1997; Sengpiel, Baddeley, Freeman, Harrad & Blakemore, 1998), but it does capture many (e.g. Albrecht & Geisler, 1991; Geisler & Albrecht, 1992; Heeger, 1992a,b; Carandini et al., 1997; Nestares & Heeger, 1997; Tolhurst & Heeger, 1997a,b) and is sufficient for our purposes here. This normalization has also been proposed in models of higher levels of the visual cortex, e.g. MT (Heeger, Simoncelli & Movshon, 1996; Simoncelli & Heeger, 1998). A normalization process may prevent overload on higher levels and overcome the limitations of a restricted dynamic range while simultaneously preserving selectivity along dimensions like orientation and spatial frequency. (See discussions and references in, e.g. Bonds, 1993; Victor, Conte & Purpura, 1997; Lennie, 1998). Recently, it has been suggested that such normalization or contrast-gain

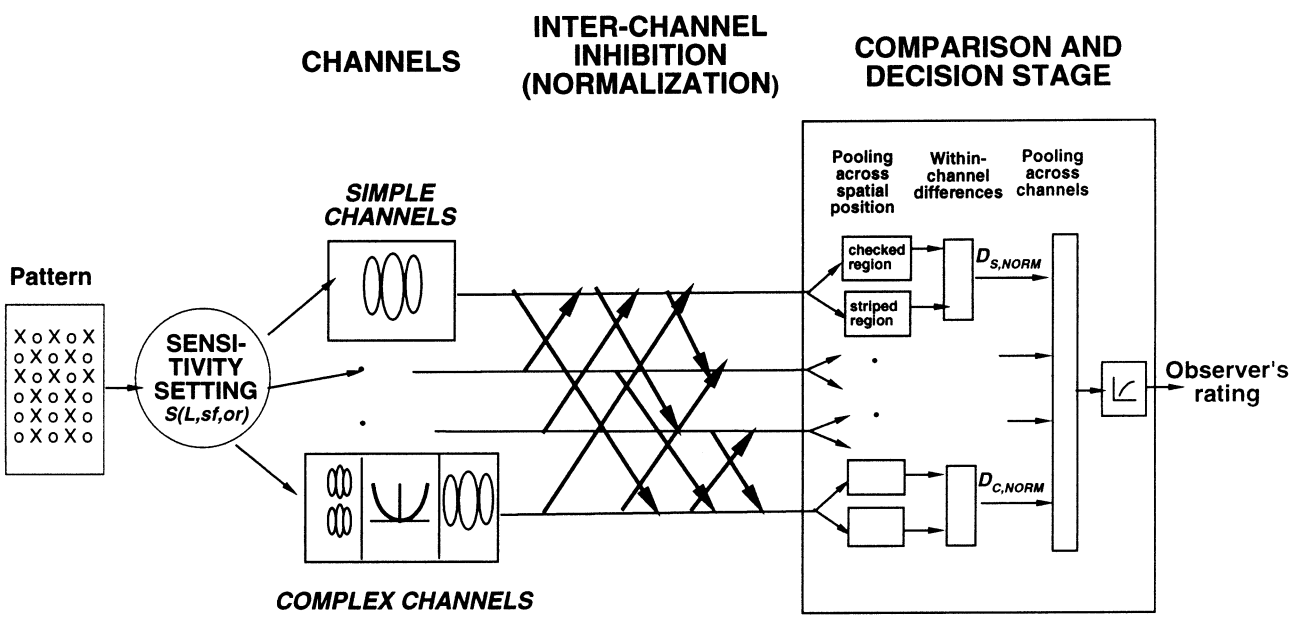

Fig. 1. The normalization model. In this model the compressive effects in texture segregation and related tasks result from inhibitory interaction among the outputs of spatial-frequency- and orientation-selective channels, as in a contrast-gain-control or normalization network. The early sensitivity-setting stage shown on the diagram does not introduce any compression for contrasts less than 100%. It does set a sensitivity factor that depends on mean luminance, spatial frequency, and orientation. It may be thought of as a linear filter for any fixed mean luminance. It also could be incorporated into the channels themselves. This model is shown here to be consistent both with the results from constant-difference experiments (here and Graham & Sutter, 1996) and the results from area experiments (Graham & Sutter, 1998).

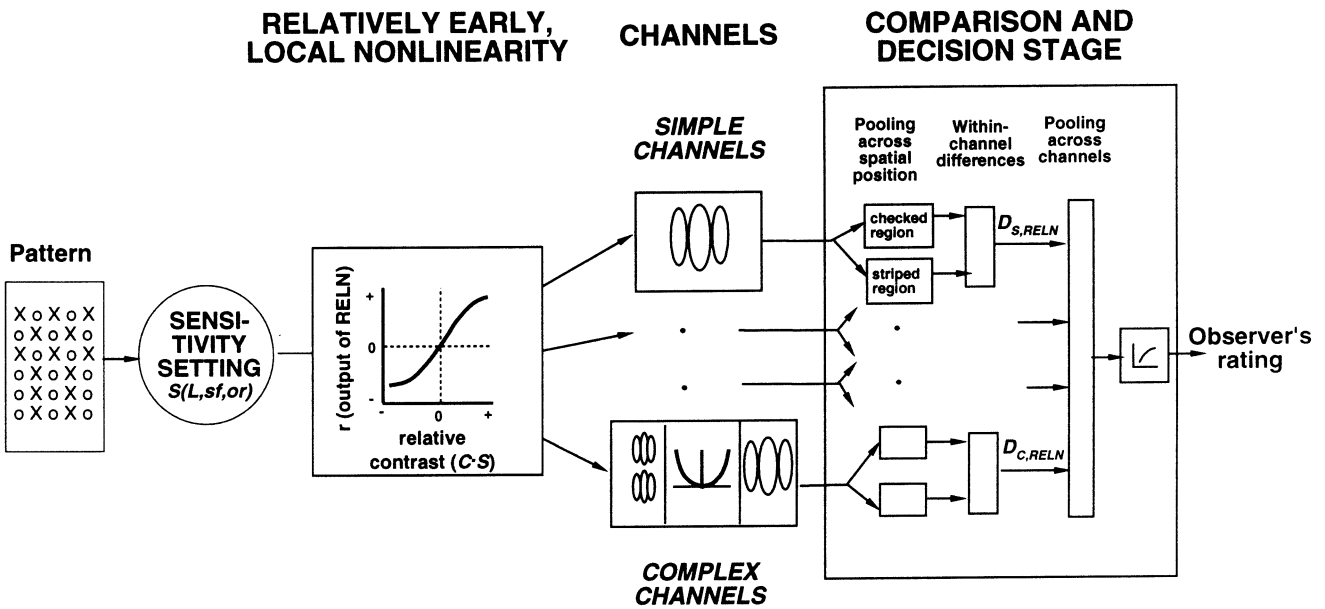

Fig. 2. The relatively-early-local model. In this model the compressive effects result from an early, local (pointwise) nonlinear function that occurs before the spatial-frequency- and orientation-selective channels. Note that there is a sensitivity-setting stage that comes before the early-local nonlinear function. (The early sensitivity-setting stage does not introduce any compression for contrasts less than 100%. It does set a sensitivity factor that depends on mean luminance, spatial frequency, and orientation.) This relatively-early-local model can explain the results of all our constant-difference experiments; however, the original form of early-local model (without an early sensitivity-setting stage) cannot (Graham & Sutter, 1996). We show here that this relatively-early-local model cannot simultaneously explain the results of constant-difference and area experiments and thus must be rejected.

control has the characteristics needed to help encode natural images efficiently (Simoncelli & Schwartz, 1998; Zetzsche, Krieger, Schill & Treutwein, 1998). Many investigators have invoked inhibition among channels — sometimes explicitly in a normalization network — in order to explain behavioral results from texture segregation and related tasks (e.g. Malik & Perona, 1990; Graham, 1991; Graham et al., 1992; Solomon, Sperling & Chubb, 1993; Wilson & Humanski, 1993; Bowen & Wilson, 1994; Foley, 1994; Teo & Heeger, 1994; Smith & Derrington, 1996; Rohaly, Ahumada & Watson, 1997; Watson & Solomon, 1997; Wilkinson, Wilson & Ellemberg, 1997; Foley & Schwarz, 1998; Olzak & Thomas 1998; Snowden & Hammett, 1998.)

The sensitivity-setting stage in Fig. 1 resets sensitivity as a function of the spatial frequency, orientation and mean luminance. However, on a constant mean luminance, it does not introduce any nonlinearity for contrasts below 100% and thus, in these conditions, may be approximated by a linear spatial filter. It incorporates at least the lens of the eye as well as some spatial filtering and light adaptation in the retina and LGN.

All the results reported below will turn out to be consistent with a normalization network in conjunction with complex channels (having an expansive rectifying function at the intermediate stage) and simple channels (Fig. 1).

1.2.2. Early local nonlinear function

The second candidate (Fig. 2) for the intensive nonlinearity is an *early local nonlinear function*. By 'early' we mean 'before the channels', and thus it could occur either at the retina or at the LGN, perhaps as a result of processes whose primary purpose is to adapt to the light level (for a recent review of retinal light adaptation see Hood, 1998) or as a result of an early contrast-gain control (e.g. Shapley & Victor, 1978, 1979, 1981; Bernardete, Kaplan & Knight, 1992). By 'local' we mean pointwise, or, more exactly, 'dependent on an area very small relative to the receptive fields characterizing the channels'.

In a strong form of the early-local hypothesis, a nonlinear function is applied directly to each point of the stimulus. This function is compressive both for decrements and increments from the mean luminance, and its output is the input to the spatial-frequency and orientation-selective channels. This strong form of the early-local model was ruled out by the results of changing the scale of the texture patterns (Graham & Sutter, 1996).

However, the modified form that is shown in Fig. 2 could not be ruled out. In this modified form, *the relatively-early-local model*, the pointwise compressive function occurs after a sensitivity-setting stage (like that in the normalization model) but before the spatial-frequency and orientation-selective channels.

On first consideration, it might seem that any local process could be ruled out by previous results showing that judgments depending on single separated texture elements are apparently much less compressive than judgments about regions involving exactly the same elements (Beck, Graham & Sutter, 1991). But, as discussed further in Graham and Sutter (1996), there is a possible alternative interpretation of these results, and they do not rule out the relatively-early-local model.

One of the main conclusions from the studies below will be to rule out definitively even the *relatively*-early-local model of the intensive nonlinearity in texture segregation.

1.3. This study

This paper presents some new results of *constant-difference experiments* and compares these results (as well as some previously-published results from another kind of experiment) to predictions from the models described above. *Constant-difference experiments* were introduced in Graham (1991) and Graham et al. (1992) and have been very useful in investigating both the spatial and the intensive nonlinear processes. The experiments here use element-arrangement patterns, some examples of which are shown in Figs. 3 and 4. In element-arrangement patterns, there are two types of elements, and the two texture regions differ in how these elements are arranged (striped in the side regions and checkerboard in the center region in Figs. 3 and 4). In each stimulus used here, the two kinds of elements were identical in everything except in contrast.

Square-element textures like those in panel (a) and panel (b) of Fig. 3, where the squares and between-square spaces are the same size, should be processed primarily by simple channels (Graham, 1991; Graham

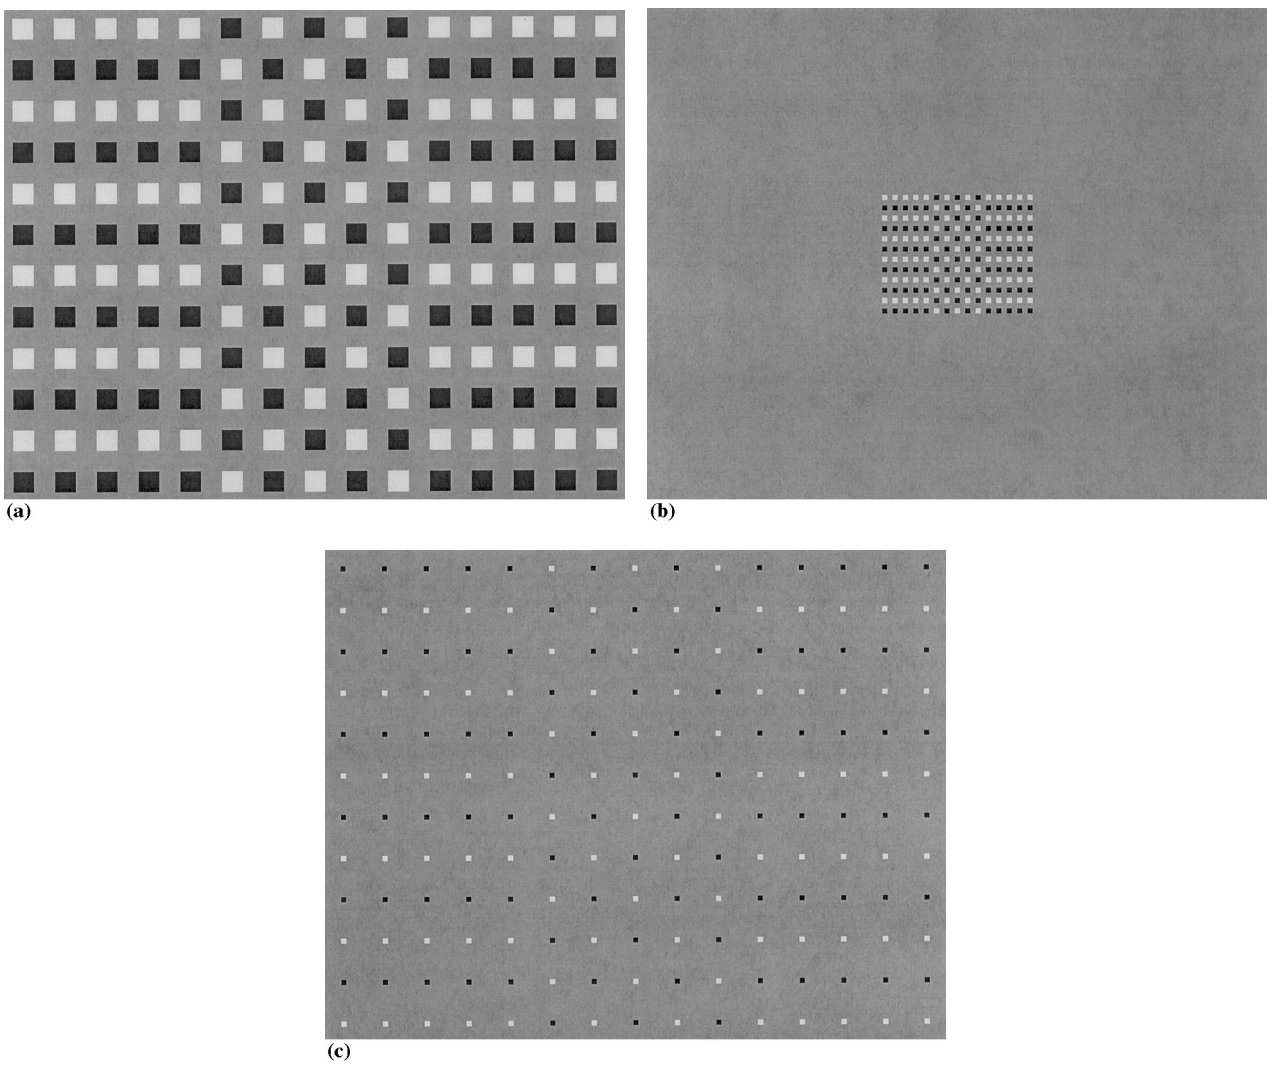

Fig. 3. Three examples of element-arrangement texture stimuli composed of square elements. Each stimulus contains three regions composed of the same two element types but distinguished by the arrangement of the elements: striped in the side regions and checkerboard in the middle region. Panel (a) shows a *large*/*regular* pattern, panel (b) a *small*/*regular* pattern, and panel (c) a *small*/*sparse* pattern. These are all *opposite-sign-of-contrast* stimuli (*contrast-ratio angle* of 0°).

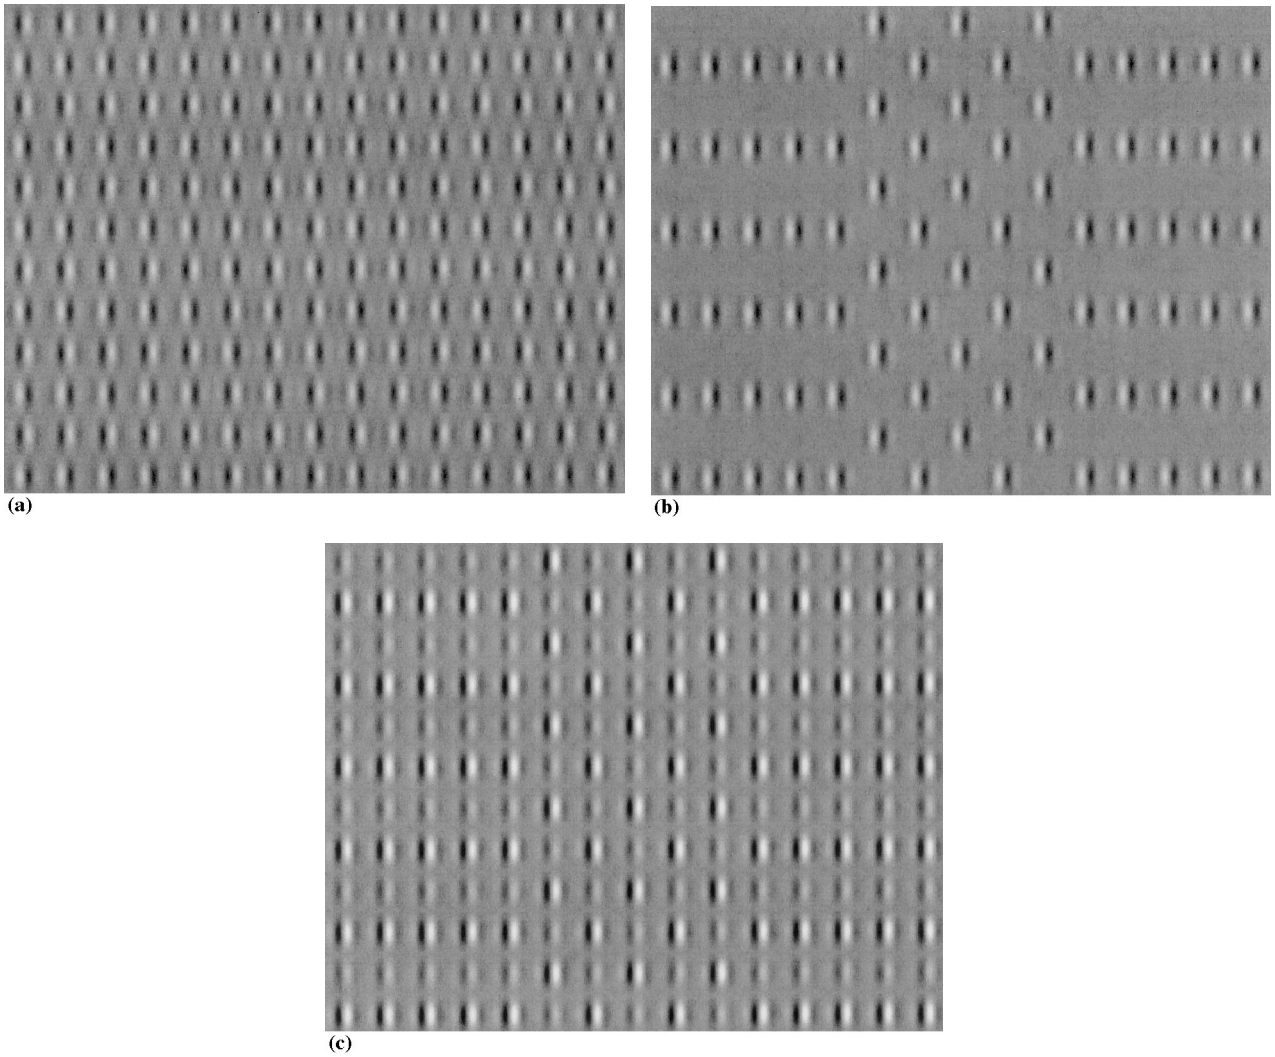

Fig. 4. Three examples of the element-arrangement texture stimuli composed of grating-patch elements: *opposite-sign-of-contrast* (panel (a)), *one-element-only* (panel (b)), *same-sign-of-contrast* (panel (c)). This figure shows the lower spatial-frequency grating elements used in this study (3 c/deg as viewed by the observers). The higher-spatial-frequency grating elements (12 c/deg) were characterized by a period one-fourth that shown here, so they had four times as many cycles per element, but the patterns were otherwise identical to these (the element's window size and the spacing of the elements were the same). The stimuli in this figure were supposed to be from the same *constant-difference series* but reproduction will have distorted them.

et al., 1992; Graham & Sutter, 1996). We will refer to these as *regularly-spaced* square-element patterns, and most of our previous constant-difference experiments have used textures like these. Another kind of square-element pattern is also used here (panel (c) of Fig. 3): a *sparse* pattern in which the spaces between the squares are much larger than the squares. As will be shown here, complex channels may be more sensitive than simple channels to sparse patterns of square elements even though simple channels can in principle segregate them.

Here we also report extensive constant-difference experiments using grating-patch element-arrangement patterns, examples of which are shown in Fig. 4. Again the two element types in a given stimulus are identical except for contrast. Note that, for grating-patch elements switching from negative to positive contrast is equivalent to a 180° phase shift. These grating-patch element-arrangement stimuli cannot be segregated by simple channels because there is little or no energy at the frequency/orientation combinations which distinguish the checkerboard from the striped regions. Hence, within this framework they must be processed by complex channels.

For both grating-patch-element and square-element textures, this paper reports experiments over a wide contrast range including both lower and higher contrasts than in previous experiments.

Importantly, the constant-difference experiments here include conditions allowing their results to be directly compared to results of *area experiments*. In the area experiments the tradeoff between the area and the contrast of elements was measured for both grating-patch-element and square-element textures (Graham & Sutter, 1998). We were also able to use many of the same observers here, and we ran both types of experiments with both square-element and grating-patch-element textures on a new observer as well, thus permitting a very secure comparison.

The results of this study allow us to distinguish definitively between the normalization and relatively-early-local models of the intensive nonlinearity. As described below, the normalization predictions are consistent with all the results of both constant-difference and area experiments, but the predictions of the relatively-early-local model are dramatically wrong.

2. Methods and procedures

2.1. About constant-difference series and contrast-ratio angles

In a *constant-difference series* of patterns — diagrammed in Fig. 5 — the difference between the contrasts of the two element types in a pattern remains fixed (as do their spatial characteristics), but the contrasts of both element types vary together. Each *constant-difference experiment* investigates a number of constant-difference series of the same pattern. (We use the word 'pattern' here to mean a particular spatial arrangement and a particular background luminance without commitment to the contrast of the elements).

More precisely, each *constant-difference experiment* uses all distinct combinations of evenly-spaced contrast levels in the two element types of some pattern, as

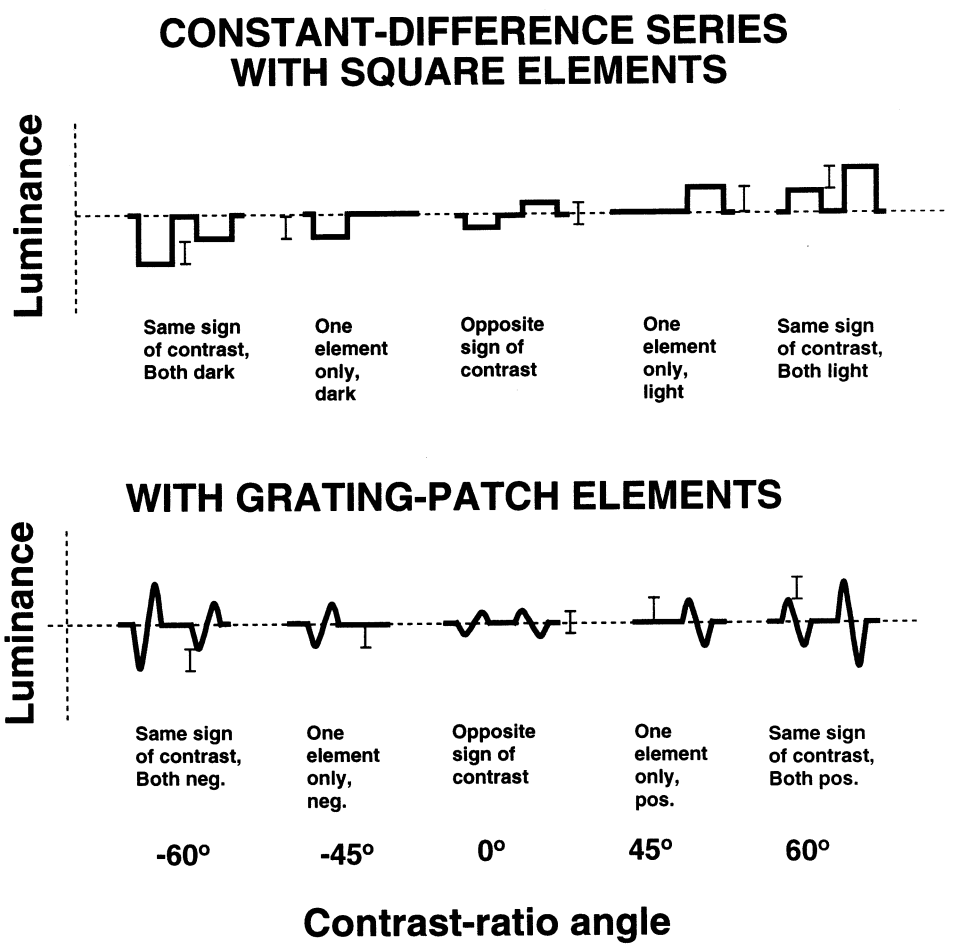

Fig. 5. Diagram illustrating two constant-difference series. Each of the five small diagrams in the top (or bottom) row shows a sketch of the luminance profile of the two element types in a square-element (or grating-element) stimulus like those in Fig. 3 (or 4). The five stimuli in each row are all in the same *constant-difference series*: in particular, the difference between the luminances (or equivalently in these cases, the contrasts) of the two element types is the same, as indicated by the small vertical bars. The contrast of the grating patch is arbitrarily taken to be negative when the bar to the left of center is dark and positive when it is bright. The sketches here show each element as containing only one cycle of sine-wave, but in the experiment each element contained more. The *contrast-ratio angle* labeling the bottom of the figure is a convenient measure for keeping track of the members of a series and is illustrated further in the next figure.

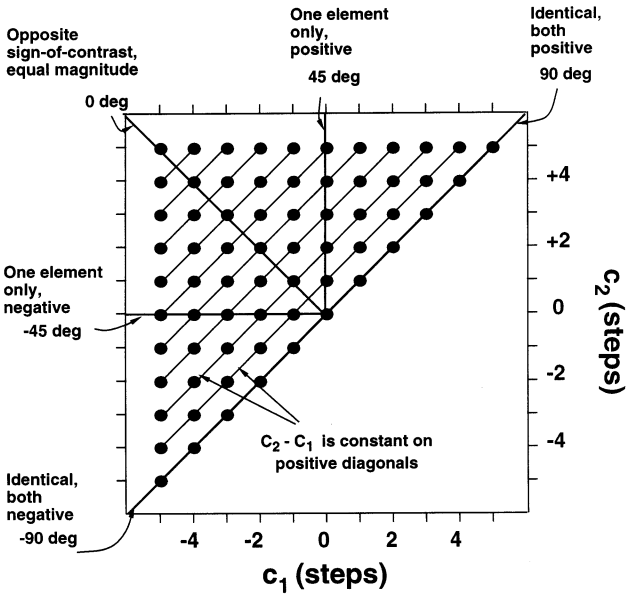

Fig. 6. Diagram of the full set of stimuli used for each *constant-difference experiment*. The contrast of one element type is plotted on the horizontal axis and that of the other element type on the vertical axis. (The contrasts are in arbitrary steps. For the value of the step in each experiment, see text.) The ratio of the contrasts of the two element types is constant along lines through the origin. The corresponding *contrast-ratio angle* is labeled outside the diagram. The stimuli along each positive oblique line form a *constant-difference series*. Each experiment with square-element textures consisted of the 66 stimuli as shown here — that is, with ± 5 steps of contrast (producing 11 different contrast levels). Each experiment with grating-element textures consisted of 91 stimuli formed by using ± 6 contrast steps (producing 13 different contrast levels) There were 11 different series in each square-element experiment (corresponding to the 11 diagonal lines of positive slope in the diagram here) and 13 different series in each grating-element experiment. Note that each constant-difference series contains one stimulus less than the series below it (the series for the next lower difference); the series for the largest difference contains only one stimulus.

diagrammed in Fig. 6. In each of the square-element experiments there are 11 levels of contrast (as shown in Fig. 6), and in each of the grating-element experiments, there are 13 levels. The 11 (or 13) levels include zero, and are evenly-spaced and symmetric around zero. This produces 66 (or 91) different stimuli in an experiment. The contrast range is determined from the magnitude of the minimal non-zero contrast, (a 'step' in Fig. 6); the maximal positive contrast is five times the step size for square-element experiments (and six times for grating-element experiments).

Stimuli along any positive diagonal in Fig. 6 form a constant-difference-series. Thus, each square-element constant-difference experiment reported here contains 11 series (as illustrated by the positive diagonals in Fig. 6), and each grating-element experiment contains 13 series.

It is useful to characterize each stimulus by its *contrast-ratio angle*. This measure is specified in the bottom labels in Fig. 5 and the left and top labels of Fig. 6. (All stimuli on a line through the origin in Fig. 6 have the same ratio of contrasts in their two elements. Such a line can be represented by its angle, which we measure relative to the negative diagonal.) At a contrast-ratio angle of zero, the elements have opposite-but-equal contrasts. Square-element opposite-sign-of-contrast stimuli are shown in all three panels of Fig. 3. In the case of grating-patch elements, elements of opposite contrast are 180° out of phase (e.g. panel (a) of Fig. 4). Angles within 45° of zero characterize stimuli in which the two element types are of opposite sign of contrast but potentially different magnitudes of contrast. Angles equal to $+45$ or $-45°$ characterize stimuli in which the contrast of one element type has been reduced to zero, so only the other element type is visible (e.g. panel (b) of Fig. 4). Angles greater than 45° in magnitude characterize stimuli in which the two element types have the same sign of contrast (same phase for grating-patch elements) although generally different magnitude (e.g. panel (c) of Fig. 4).

2.2. Overview of constant-difference experiments

With square elements, nine experiments were done for each of six observers. Results from the nine experiments with square elements will be shown in Figs. 8 and 9 for two individual observers. The nine experiments resulted from combinations of three square size/spacing conditions (illustrated in Fig. 3 and described below — corresponding to the three rows in Figs. 8 and 9) with three contrast ranges (step sizes of 1.33, 4 and 12% contrast — corresponding to the three columns in Figs. 8 and 9). Trials from the nine square-element experiments for a given observer were randomly intermixed. Each session consisted of one trial of each stimulus in each experiment. Four sessions were run for each observer. Different sessions were usually run on different days.

With grating-patch elements, six experiments were done for each of four observers. Results from the six experiments with grating-patch elements will be shown in Figs. 10 and 11 for two individual observers. The six experiments resulted from combinations of two spatial frequencies in the Gabor patches (3 c/deg and 12 c/deg — corresponding to the two rows in Figs. 10 and 11) with three contrast ranges (corresponding to the three columns in Figs. 10 and 11). The three contrast ranges were determined by a step size of 1.33, 4 and 12% contrast for three of the four observers. For subject CV, who was run first, the three contrast ranges for 3 c/deg patches were determined by a step size of 1.33, 2.67 and 3% contrast, while those for 12 c/deg patches were determined by a step size of 5.33, 10.57 and 16% contrast. Trials from the six grating-element experi-

ments for a given observer were randomly intermixed. Each session consisted of one trial of each stimulus in each experiment. Four sessions were run for each observer. Different sessions were usually run on different days.

2.2.1. More details of square-element stimuli

The number, spacing, and arrangements of square elements are shown in Fig. 3.

The width of a square element was 0.33° (16 pixels at the viewing distance of 0.91 m) or 0.08° (4 pixels). For the large (0.33°) squares, the center-to-center spacing between neighboring elements was 0.67° (32 pixels), which is the same center-to-center spacing used for the grating-patch element patterns. This produces a pattern in which large squares were separated by equally large spaces (that is, a duty cycle of 0.5). We refer to these patterns as *large/regular* patterns; the opposite-sign-of-contrast case is illustrated in panel (a) of Fig. 3. Notice that the repetition period (within a given region) was two rows and two columns of elements and inter-element spaces ($1.33 \times 1.33°$ or 64×64 pixels). Thus, for these large squares, the fundamental frequency (the reciprocal of the repetition period) of either the checkerboard or the striped region was 0.75 c/deg both horizontally and vertically.

For the small 0.08° squares, there were two spacing conditions. One was a center-to-center spacing of 0.16°, leading to a pattern in which small squares were separated by equal-width small spaces, that is, a duty cycle of 0.5. This corresponded to a fundamental frequency of 3 c/deg in both horizontal and vertical directions. We will refer to these patterns as *small/regular* patterns; the opposite-sign-of-contrast case is illustrated in Fig. 3 panel (b).

The second spacing condition for the small 0.08° squares was a center-to-center spacing of 0.67° (which is the spacing used with the bigger squares), producing a pattern in which small squares were separated by much bigger spaces for a duty cycle of 1/8, or, in other words, where the density of squares was lower than in the regular cases. The fundamental frequency of these patterns was 0.75 c/deg as in the large/regular patterns. We refer to these patterns as *small/sparse* patterns; the opposite-sign-of-contrast case is illustrated in Fig. 3 panel (c).

The square-element patterns can be segregated by simple linear channels tuned to the fundamental frequency (0.75 c/deg for large/regular and small/sparse, and 3 c/deg for small/regular patterns). They could also in principle be segregated by complex channels with first-stage filters tuned to high spatial frequencies (present in, e.g. the edges of the squares) and second-stage filters tuned to the fundamental frequency if these complex channels are sensitive enough. On the basis of previous experiments done with *regular spacing*, one expects these patterns to be segregated primarily by simple channels although with some minor complex-channel contribution. The question of which channels segregate the sparse-spacing condition is investigated here.

2.2.1.1. Relationship to stimuli from previous area experiments. The *large/regular* and *small/regular* patterns used here are just like the 0.33 and 0.08° square patterns used in constant-difference experiments by Graham and Sutter (1996). Also, the large/*regular* and *small/sparse* patterns here have elements spaced like those in the square-element area experiments of Graham and Sutter (1998) and the sizes of the square elements here correspond, respectively, to the 0.33 and 0.08° square elements in those area experiments.

2.2.2. More details of grating-patch-element stimuli

The number and arrangements of elements were identical to that for square elements and are shown in Fig. 4.

The width of a Gabor-patch element (full width at half peak) was always 16 pixels (0.33°). The harmonic oscillation was in sine phase with respect to the window. The orientations of all the Gabor patches were vertical. Their spatial frequency was either 3 c/deg (a period of 16 pixels) or 12 c/deg (a period of 4 pixels).

The center-to-center spacing between neighboring elements was 32 pixels (0.67° at the viewing distance of 0.91 m). The repetition period (within a given region) was two rows and two columns of elements and was $1.33 \times 1.33°$ (64×64 pixels). Thus the fundamental frequency (the reciprocal of the repetition period) of either the checkerboard or the striped region was 0.75 c/deg both horizontally and vertically.

These grating-patch-element patterns *cannot* be segregated by a model containing only simple (linear) channels. Although some linear filters do respond differently to the checkerboard versus striped arrangements, the magnitude of these differences (as processed by any of the broad family of decision and pooling stages we consider in our models) is extremely small. Indeed, the predictions for grating-element patterns from a model composed only of simple channels (plus our usual decision and comparison stage) are very much like those that we have previously published for center-surround element patterns (Graham et al., 1992, Fig. 7).

These grating-patch-element patterns can, however, be segregated by complex channels, in particular by the complex channels with first-stage filters tuned to 3 or 12 c/deg (corresponding to the spatial frequency of the patches) and second-stage filter tuned to 0.75 c/deg (the fundamental frequency).

2.2.2.1. Relationship to stimuli from previous area experiments. The 12 c/deg grating-element patterns used here

correspond to patterns used in the grating-element area experiments of Graham and Sutter (1998), in particular, to those with the largest grating-patch elements.

2.2.3. Details in common for all experiments

The mean luminance was 18 ft-L.

The observer initiated a trial by pushing a key. The patterns were presented for 1 s, with abrupt onsets and offsets. After stimulus offset, a 1-s delay occurred and then a beep signaled the observer to make a response. The observer's response was a rating of the degree to which the three regions in the patterns perceptually segregated.

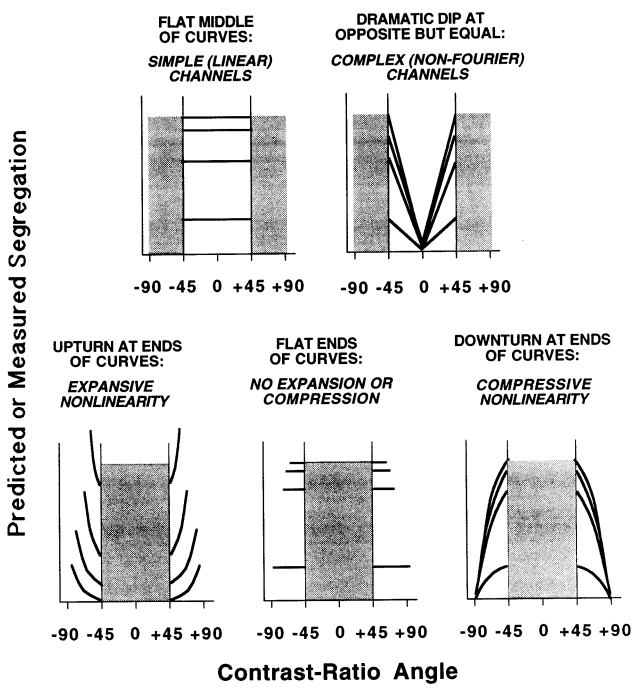

Fig. 7. Schematic illustration showing how results from constant-difference experiments can be related to the underlying model processes. The vertical axis shows segregation, either predicted or measured. The horizontal axis shows contrast-ratio angle. The vertical dotted lines at $\pm 45°$ mark the one-element-only patterns in the constant-difference series. Each curve represents the results from a constant-difference series of stimuli. Gray shading covers regions that are not particularly informative about the issue in question. As shown in the top row, points in the middle of the curves (opposite-sign-of-contrast stimuli, having contrast-ratio angles between $+45$ and $-45°$) are most useful in telling whether simple or complex channels or both are segregating the patterns. As shown in the bottom row, points at the ends of the curves (same-sign-of-contrast stimuli, having contrast-ratio angles < -45 or $> +45°$) are most useful in distinguishing between compressive and expansive nonlinearities. Such expansive or compressive nonlinearities may occur at several places in the model, in particular, at the intermediate stage of the complex channels or as an early-local nonlinear function, or as the result of a normalization network. (In the bottom panels, each curve extends over a slightly narrower range of contrast-ratio angles than the curve below it because, in our experiments, each series extended over a narrower range of contrast-ratio-angles than series characterized by lower differences).

Viewing was binocular at a distance of 0.91 m.

(Further details can be found in Graham & Sutter 1996, 1998).

2.2.4. Observers

There were six observers in the square-element experiments reported here. Four of these also ran in the grating-element experiments. All of them were between 17 and 35 years of age and had normal or corrected-to-normal acuity.

2.3. Calculating model predictions and fitting the experimental results

In this study we approximate models' predictions by simple equations, as we have done in some of our earlier papers. Appendix A here summarizes these equations from previous papers. These equations treat each element type as an entity and perform simple algebra on element types (rather than explicitly calculating the spatial filterings and other transformations at each stage in the model). Partly as a consequence of the physical separation of the elements in our patterns, the rather simple equations are a very good approximation to the results of the full model. (Earlier papers have shown some examples of predictions from the full model and also some comparisons of these full predictions to predictions from approximate equations: Sutter et al., 1989; Graham, 1991; Graham et al., 1992). Not only does this approximate-equations approach make the computational task substantially easier but also, importantly, the simplicity of the equations makes it easier to see why the models make the predictions they do.

This approximate-equations approach does, of course, introduce approximations that then need to be carefully considered for possible import. There are two separate types of approximation inherent in this approach. First, in our equations we frequently concentrate only on the 'tuned' channel or channels (those most able to segregate the patterns under consideration). There are undoubtedly less-well-tuned channels also, and consideration of these less-well-tuned channels could introduce subtleties like those discussed at length for area experiments in Graham and Sutter (1998). However, calculations exploring the possible intrusions in the experiments here (predictions from model versions containing realistic less-well-tuned channels) suggest that they play only a minor role here and introduce no serious artifacts into the conclusions. (Indeed, the intrusion of these other channels seems to be less of a problem in constant-difference experiments, where the elements in any one stimulus are all the same size, than in the area experiments.) Secondly, but of lesser possible significance in these experiments, the approximate-equations approach neglects extra frequencies and orientations introduced into the processing by the pointwise nonlinearities (i.e. the early, local

function and the function at the intermediate stage of the complex channels).

The conclusions stated below are based on calculating predictions from many versions of the relatively-early-local and normalization models. We varied the properties of the normalization network, and/or the relatively-early-local nonlinearity, and we also varied the exponent on the function at the intermediate stage of the complex channel. We also did one large set of quantitative fits of models to the results of these experiments, analogous to those reported in Graham and Sutter (1996). See Appendix C for more information.

3. Results and discussion

Experimental results and model predictions for the constant-difference experiments will be plotted in a form which is useful for distinguishing among models. This form is illustrated in each panel of Fig. 7. The horizontal axis shows the *contrast-ratio angle*. The vertical axis shows *segregation* — either measured or predicted. (The observer's ratings of perceived segregation will not, in general, be directly proportional to predicted segregation, but the two are assumed to be related by some monotonic transformation, as illustrated in Fig. 16 of Graham et al., 1992.) Each curve in these plots connects results for the stimuli from one constant-difference series. Fig. 7 shows several possible characteristics of results plotted in this form and how they correspond to underlying processes in the models. Although both the spatial and intensive nonlinearity can affect all parts of a constant-difference curve, the magnitude of their effects is different in different places along the curves, and thus it is useful to divide the curves in the plots into middles (see top panels Fig. 7) and ends (see bottom panels Fig. 7) and consider what results in each region correspond to what processes in the models. Since previous papers have described these effects in detail, we just briefly survey them here.

As indicated in the top row of Fig. 7, the middle region of the curves (opposite-sign-of-contrast stimuli, contrast-ratio-angles between -45 and $+45°$) is particularly useful for deciding whether simple channels or complex channels (or both) are segregating the patterns. Simple channels lead to flat curves since all stimuli in the same constant-difference-series are approximately equally segregatable by simple channels. (See Fig. 7 of Graham et al., 1992, for a demonstration of this.) However, complex channels lead to a dramatic dip in segregatability at the opposite-but-equal case because the second stage of the complex channels cannot tell the difference between equal-but-opposite elements once they have passed through the first stage and the rectification-type function at the intermediate stage. (This is illustrated for square elements in Fig. 11 of Graham et al., 1992; an analogous illustration could be drawn for grating elements).

As indicated in the bottom row of Fig. 7, the ends of the curves (same-sign-of-contrast patterns, contrast-ratio-angles less than $-45°$ or greater than $+45°$), are useful for deciding whether there are expansive or compressive nonlinearities acting in the system. Upturned ends of the curves indicate expansion; conversely, downturned ends indicate compression. With simple channels and with complex channels having a piecewise-linear function (e.g. conventional full-wave or half-wave rectification) at the intermediate stage, the ends of the curves directly reflect the nature of the intensive nonlinearity in the models. Indeed it was the typical downturn in experimental results that initially led us to posit a compressive intensive nonlinearity. But, as we discuss below, the ends of the curves may also reflect the intermediate stage in complex channels.

Figs. 8 and 9 show the full results for two observers from the nine constant-difference experiments with square elements. Figs. 10 and 11 show the full results for those same two observers from the six constant-difference experiments with grating elements. In all four figures, the three columns show the results from the three different contrast ranges. The three rows in Figs. 8 and 9 correspond to the three kinds of square-element patterns (different sizes and spacings of squares). The two rows in Figs. 10 and 11 correspond to the two spatial frequencies of grating-patch elements. Each data point shows the results for one stimulus. These two observers' results look generally similar, although they differ somewhat from each other. These two observers together are quite representative of the results from the other observers not shown here. We will discuss these results at some length in the next several subsections.

3.1. Which patterns are segregated by simple channels, and which by complex channels?

First note that the middle regions of the curves in Figs. 8–11 are consistent with the suppositions we have made before that: (i) the grating-patch element patterns are segregated by complex channels while; (ii) the regularly-spaced square-element patterns are primarily segregated by simple channels (although with some complex-channel intrusion). There is a new and different result here, however, for the sparsely-spaced square-element patterns. These patterns show a much greater amount of complex-channel intrusion than do regularly-spaced squares. In fact, for three of the six observers (WS, who is shown here in Fig. 9, CAS and CH), perceived segregation descends to zero for equal but opposite contrasts, thus producing plots with a double-peaked appearance just like that for the grating-patch element results (e.g. Figs. 10 and 11). This double-peaked appearance is the result expected if complex channels are very heavily involved in the segregation of

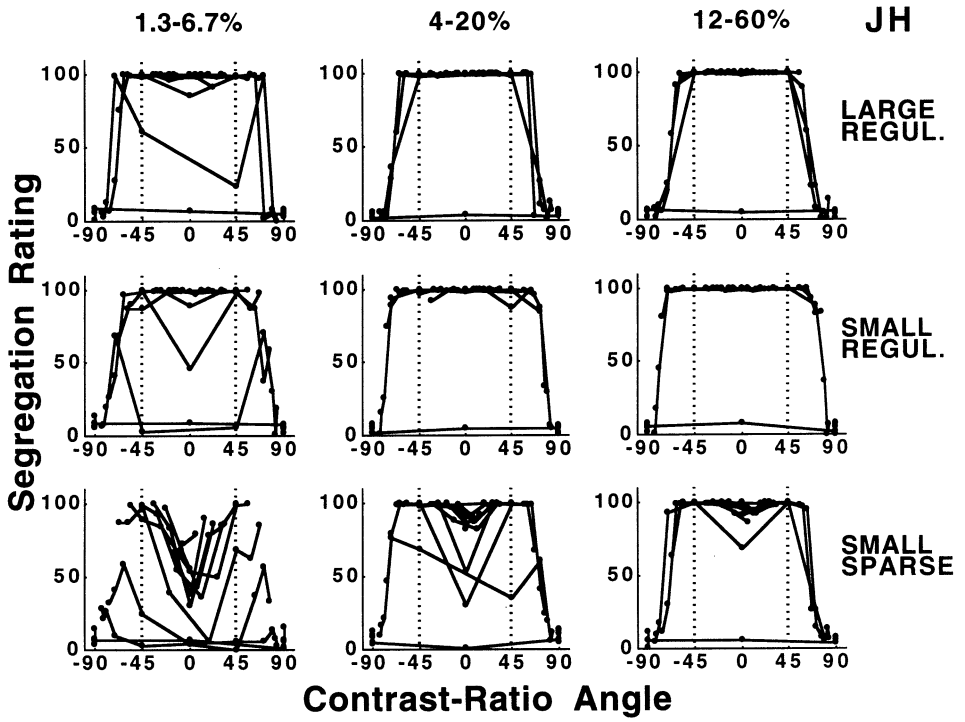

Fig. 8. Results from the square-element constant-difference experiments for one observer (JH). Experiments were done with three kinds of square-element patterns (three rows) and for three contrast ranges (three columns). In each panel, perceived segregation is plotted on the vertical axis as a function of contrast-ratio-angle on the horizontal axis. The vertical dotted lines mark the one-element-only patterns in the constant-difference series. Each point represents one of the 66 stimuli studied in an experiment and shows the segregation rating averaged over four presentations of that stimulus done in four separate sessions. The lines connect stimuli in a particular constant-difference series.

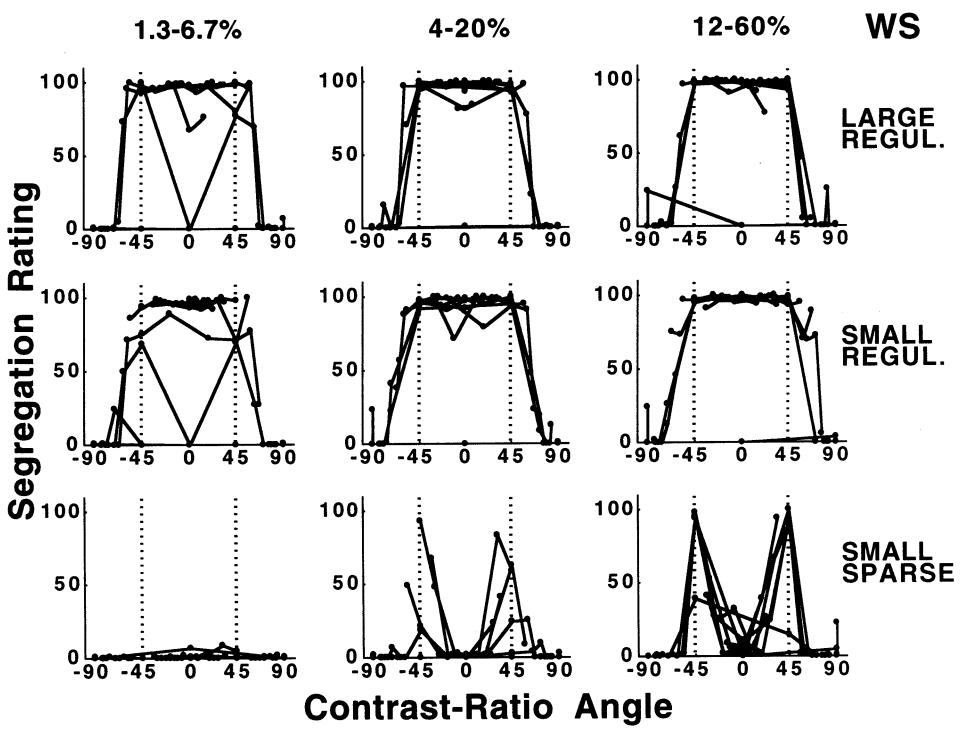

Fig. 9. Results from observer WS for square-element patterns in same format as Fig. 8.

these sparse-spacing patterns. Such heavy involvement is perfectly plausible as the ratio of complex-to-simple channel sensitivity should increase substantially as one goes from regular to sparse spacing (see Appendix B for more explanation). To look at the matter from a different perspective, as small elements get further apart (a comparison of *small-regular* versus *small-sparse*) complex channels will generally become better suited than simple channels for tying those elements together into larger patterns. Substantial individual differences in the amount of complex channel are seen in these results as they were in the area experiments (Graham & Sutter, 1998) and may be due to the same cause: individual differences in the sensitivities of certain complex and simple channels.

3.2. Failure of the early-local model

Now note that the ends of the measured constant-difference curves (same-sign-of-contrast patterns) turn down in almost all cases This downturn is seen clearly in all the panels of Figs. 8–11 except for the lower left panels. The lower left panels sometimes even show an upturn (which is particularly clear in the lower left panels of Figs. 8 and 11). This upturn at very low contrasts for certain stimuli will be discussed later when it can be interpreted in the light of conclusions from the other results.

A compressive early local model — either of the original strong or the modified relative form — easily

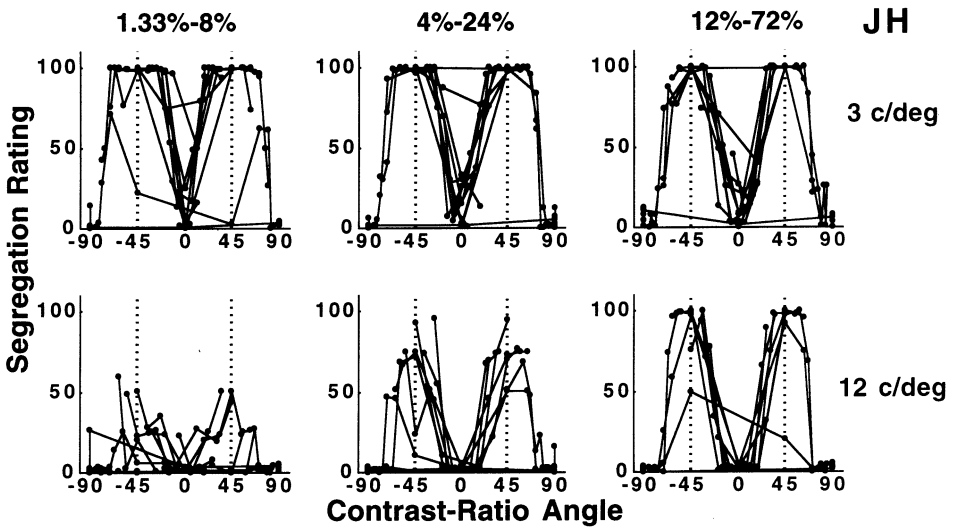

Fig. 10. Results from grating-patch-element constant-difference experiments for one observer (JH). Experiments were done with two different grating-patch spatial frequencies (two rows) and three different ranges of contrast (three columns). Format of the individual panels is like that in Figs. 8 and 9. Each point represents one of the 91 stimuli studied in an experiment and shows the segregation rating averaged over four presentations of that stimulus done in four separate sessions.

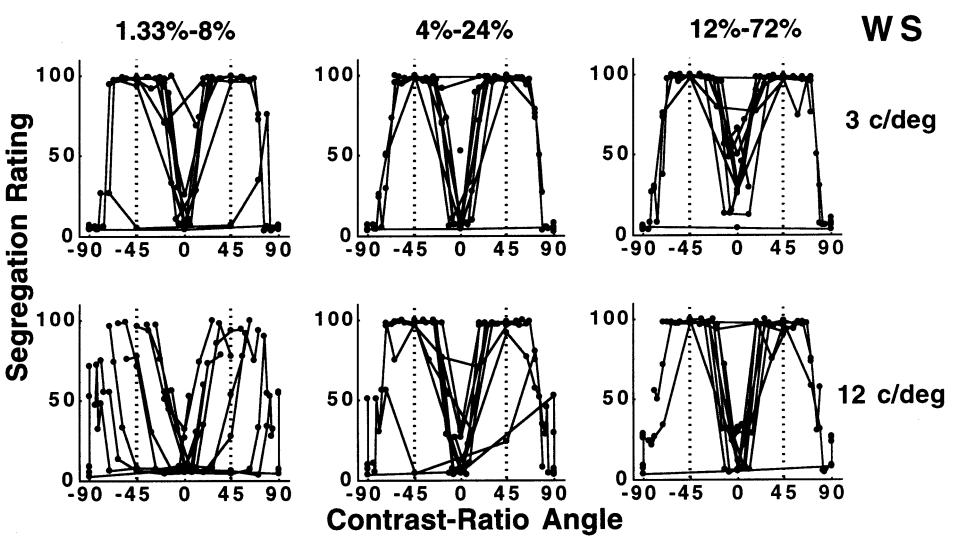

Fig. 11. Results from observer WS for grating-patch element experiments in same format as Fig. 10.

accounts for the decreased segregatability toward the ends of the curves in constant-difference results. This decreased segregatability occurs because the compressive early-local nonlinear function is centered at the background luminance (like that shown in Fig. 2), thus protecting discriminability among luminances near the background luminance while sacrificing discriminability further away. A diagram illustrating this logic can be found in Fig. 18.8 of Graham (1991) or Fig. 15 of Graham et al. (1992).

However, as discussed next, the early-local model (either of the original or modified form) can not predict simultaneously both the constant-difference and the area results.

3.2.1. Prediction of the early-local model

The early-local model, either of the original or the modified form (Fig. 2), turns out to make a strong prediction for the effects of area experiments and constant-difference experiments considered together, namely:

> *If compression (or, respectively, expansion) shows up in the results of one of these kinds of experiments, then compression (or, respectively, expansion) must also show up in the results of the other kind of experiment (when using the same patterns, at the same contrasts, with the same observers).*

(For a description of how compression and expansion show up in the results of area experiments, see Graham & Sutter, 1998).

To provide some insight into this prediction of the early-local model, let's go through the case of simple and complex channels separately.

3.2.1.1. Simple channels. Consider the relatively-early-local model of Fig. 2 and consider the case of regularly-spaced square-element patterns, which are processed primarily by simple channels. In this case the degree of compression or expansion in both types of experiments will be determined by the relatively early, local nonlinear function itself (see Fig. 2 and r in Eqs. (4) and (5) of Appendix A). There are no other processes in the model which can introduce compression or expansion in the simple-channel pathway in these experiments. (The sensitivity-setting stage is linear for all patterns at a given mean luminance, that is, it does not introduce either compression or expansion for the patterns used in any one of these experiments; and the comparison and decision rule stage does not introduce compression or expansion either.) Since the relatively-early-local nonlinear function remains the same in both kinds of experiments, both kinds should show compression or both should show expansion. (The original form of the model is like Fig. 2 but missing the early sensitivity-setting stage, so the prediction holds for it as well).

3.2.1.2. Complex channels. Consider the relatively-early-local model of Fig. 2 and the case of grating-patch elements which are processed by complex channels. Here the degree of compression or expansion will be determined by the combined effects of the relatively early, local nonlinear function and the pointwise function at the complex channels' intermediate stage. While there is linear spatial filtering before and between those two pointwise nonlinearities, they act in this regard approximately as a single concatenated function. For example, if the relatively-early-local function were a (compressive) power function with an exponent of 0.5 and the function at the intermediate stage of the complex channels were an (expansive) function with an exponent of 3, then overall the experimental results would look expansive with an exponent of about $3 \times 0.5 = 1.5$. And this would be true for both area experiments and constant-difference experiments.

3.2.2. The prediction does not agree with the experimental results

To help in considering the results from these two kinds of experiments simultaneously, we fit the relatively-early-local model to the results of the constant-difference experiments reported here in order to find the relatively-early-local function which produced the best fit of the model to the results. The intermediate function in the complex channels used in these fits was assumed to be piecewise-linear so that all the compression or expansion would show up in the relatively-early-local function itself. The open circles joined by solid lines in Figs. 12 and 13 show the best-fitting relatively-early-local functions obtained for some of the experimental results. (In terms of the equations, these figures plot the best-fitting value of $r(S_1 \cdot C_1)$ as a function of C_1, see Eqs. (7) and (8) of Appendix A.) Fig. 12 shows the functions for the *large/regular* square-element results for all six observers. These patterns were identical to the largest-element patterns in the square-element area experiments (Graham & Sutter, 1998). Fig. 13 shows the functions for the 12 c/deg grating-patch elements for all four observers. These patterns were identical to the largest-element patterns in the grating-element area experiments (Graham & Sutter, 1998). Remember that all the other experimental conditions, as well as the observers, were the same in the constant-difference and area experiments. The fits for the three different contrast ranges for each observer are pinned at the points of overlap and the highest output set equal to 100. Only the positive half of the function is shown as it was assumed to be anti-symmetric in these fits. (See Graham & Sutter 1996 and Appendix C here for more details of these fits.)

The straight oblique lines in Figs. 12 and 13 are plotted for comparison purposes; the solid and dashed oblique lines have slopes of 1.0 and 0.5, respectively,

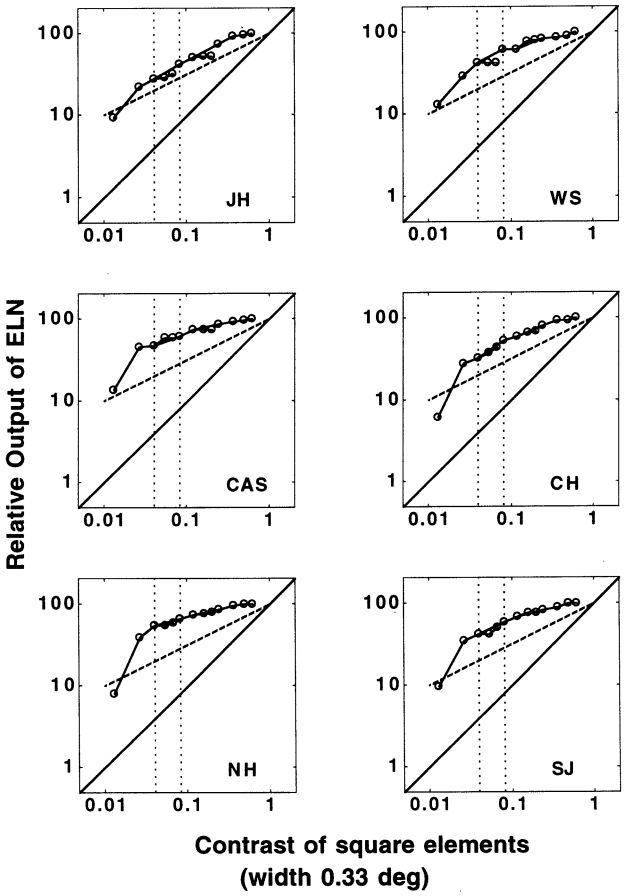

Fig. 12. The best-fitting relatively-early-local function r for square-element patterns for each of the six observers. The points in each panel show the function r for an individual observer that leads to the best fit of the relatively-early-local model to the square-element constant-difference results. (In terms of the equations, this figure plots the best-fitting value of $r(S_1 \cdot C_1)$ in Eqs. (7) and (8) of the appendix as a function of C_1.) The fits for the three different contrast ranges for each observer are pinned at the points of overlap and the highest output set equal to 100. The straight oblique lines are plotted for comparison purposes; the solid and dashed oblique lines have slopes of 1.0 and 0.5, respectively, showing a linear function and a power function with power 0.5. The vertical dotted lines show the contrast range from the area experiments run with the same stimuli in the same conditions and for the same observers (Graham & Sutter, 1998, except that CAS was run later.) Between the vertical lines the best-fitting early-local function is approximately parallel to the dashed oblique line and thus well described as a power function with a power of 0.5.

showing a linear function and a power function with power 0.5. The dotted vertical lines in Figs. 12 and 13 indicate the operative contrast ranges from the area experiments for the same stimuli. In particular, they indicate the range of contrasts over which minima in the segregation curves were measured, and it is these minima that indicate expansiveness or compressiveness (as explained in Graham & Sutter, 1998). Note that, between these vertical lines, the best-fitting relatively-early-local functions in Figs. 12 and 13 are approximately parallel to the dashed oblique line and thus well described as a power-function with a power of 0.5 (which is compressive and very close to a logarithmic function in fact). That is, the results from constant-difference experiments show compression within the ranges of contrasts used in the area experiments. The area experiments do *not* show compression within this same range. The area experiments with grating-patch elements show considerable expansion (as in a power function with an exponent of 3 or 4 typically). The area experiments with square elements show approximate linearity (where deviations from linearity are, if anything, toward expansion rather than compression). This is true for all observers in the area experiments. (Five of the observers used here are published in Graham & Sutter, 1998. Observer CAS here was run in the area experiments later and showed at least as much expansion as the published observers.)

Since (for the same range of contrasts, conditions, etc.) one experiment shows compression and the other shows expansion or linearity, the early-local model (of either form) cannot simultaneously predict the results from both constant-difference and area experiments. Thus, the early-local model of either form can be rejected as an explanation for these texture-segregation experiments.

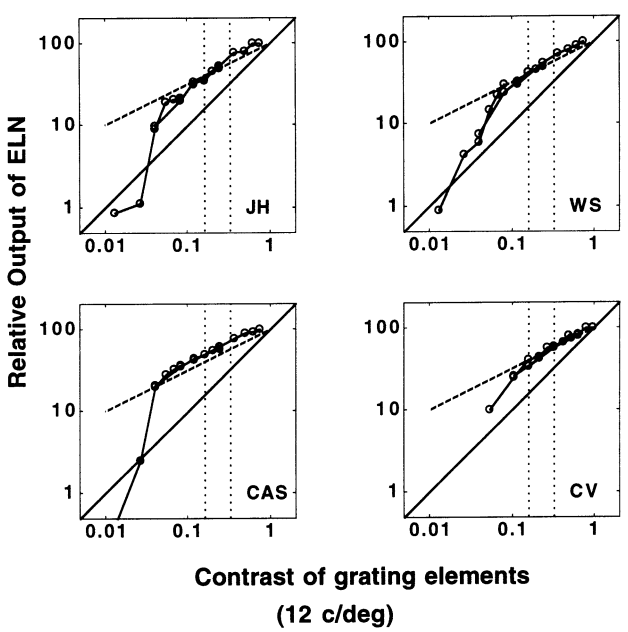

Fig. 13. The best-fitting relatively-early-local function r for grating-element patterns for each of the 4 observers. Figure in the same format as Fig. 12. Again the vertical dotted lines show the contrast range from the corresponding area experiments. Between those lines the best-fitting early-local function is approximately parallel to the dashed oblique line and thus well described as a power function with a power of 0.5.

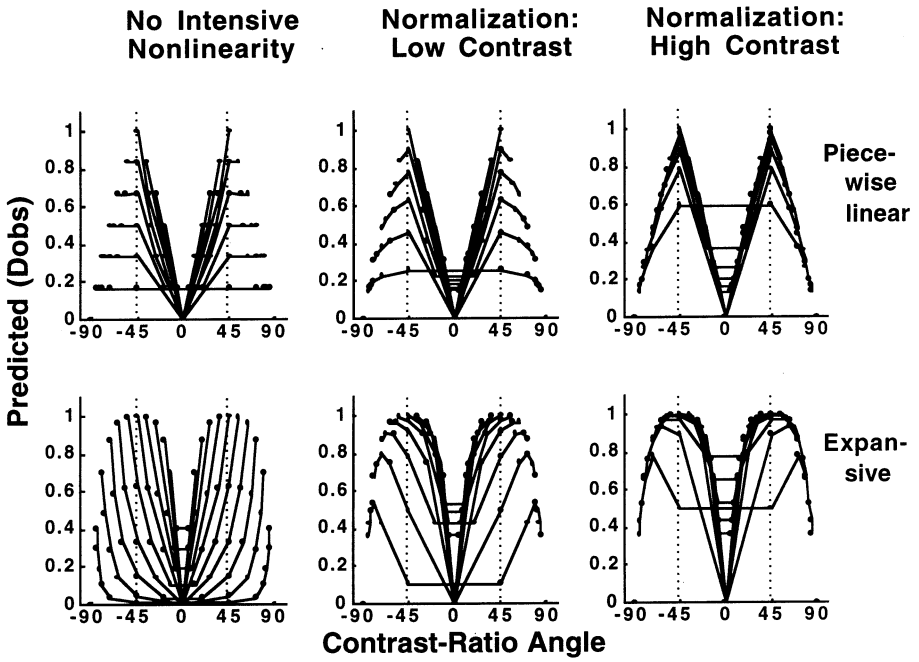

Fig. 14. Predicted segregation for constant-difference experiments when the function at the intermediate stage of the complex-channel is either piecewise linear, $k_m = 1$ (top row) or expansive with an exponent $k_m = 3$ (bottom row). The left column represents the case of no normalization (and is also the case at *very* low contrasts even if there is normalization in the model). The middle column is for contrasts low enough that the compressive effect of the normalization is still moderate. The right column is for high contrasts. The difference between the three columns was achieved by having σ in Eq. (17) of Appendix A change from very high (1000, to produce linear behavior on the left) to moderate (9, in middle column), to very small (1, in right column) while keeping the contrast values fixed (in arbitrary units) at values of $\pm 1, 2, \ldots 6$. The values of the sensitivity parameters were: $w_S = 0$, $w_X = 1$, $w_{OS} = w_{OX} = 4$. For further information see Appendix C. The predicted segregation was arbitrarily rescaled to reach a maximum of 1.0 in each panel. An observer's ratings are assumed to be a monotonic transformation of the predicted value of plotted on the vertical axis.

3.3. Successes of a normalization model

3.3.1. High-contrast compressiveness in constant-difference experiments

According to a normalization model, the decreased segregatability for same-sign-of-contrast patterns in constant-difference experiments at high contrasts (right panels Figs. 8–11, ends of curves) is easy to explain. Once the contrast is high enough to bring normalization into play, then, as one moves toward either end of a constant-difference series, there is increased inhibition from channels in the normalization pool (but at all points in the series there is approximately constant excitation from channels able to segregate the texture). See Appendix B or Graham et al., 1992 for more explanation. Quantitative fits of the normalization model to the constant-difference results (done in the same way as the fits previously reported by Graham & Sutter, 1996) show that the predictions of the normalization model are an excellent description of the results not only from square elements, as we previously showed, but also from grating elements at middle to high contrasts. These quantitative fits lead to conclusions essentially analogous to those reported in our earlier paper for square elements and are not repeated here consequently. Note that this concordance means that the normalization network is acting in the same manner on both simple and complex channels.

Some typical quantitative predictions from the normalization model for constant-difference experiments with grating elements are shown in Fig. 14. The top three panels are from a complex channel with a piecewise-linear function at the intermediate stage ($k_m = 1$ in Eq. (2) of Appendix A). The bottom row is for a complex channel with an expansive nonlinearity ($k_m = 3$). The rightmost panels are predictions with normalization for a middle- to high-contrast range. (We will discuss the left and middle panels later.) Both for piecewise-linear (top right) and expansive (bottom right) complex channels, the predicted segregation at the ends of the curves goes down, and, in fact, the ends of the different curves juxtapose as in typical empirical results for middle to high contrasts.

The downturn at the ends of the curves is predicted even for the expansive complex-channels case(bottom row, right panel) because the compressiveness in the normalization network overwhelms the expansiveness at the intermediate stage of the complex channel. (The curves for the piecewise-linear and expansive complex channels seem to have rather different shapes. But,

since a monotonic function intervenes between the predictions and the data, a shape difference of this sort is not a stable feature of the predictions and would not be discriminable in the data). Juxtaposition is predicted at the ends of the curves because, at high enough contrasts, the computation in the normalization prediction is a division of a particular channel's response by the sum of all channels' responses, and hence it is ratio of contrasts that matters. (In other words, see Appendix A, σ in Eq. (17) becomes negligible when those responses are high).

3.3.2. Expansiveness in area experiments

The normalization model (Fig. 1) can easily predict simultaneously the compressiveness seen in constant-difference experiments just discussed and the expansiveness seen in the area experiments (at the same middle to high contrast ranges, with same observers, etc.). In short — it can predict those results which eluded the early-local models. We attempt to explain why the normalization model makes these predictions in the following paragraphs.

According to the normalization model, the effects in the area experiments depend almost entirely on the properties of the channels and are almost totally independent of the normalization network while the effects in constant-difference experiments for same-sign-of-contrast patterns (at high contrasts) depend almost entirely on the properties of the normalization network and are almost totally independent of the properties of the complex channels. This near-independence predicted by the normalization model for the two types of experiments is entirely unlike the prediction above of the early-local model. The reason behind the predicted near-independence is the following: In the normalization model, the tradeoff between area and contrast occurs within individual channels (when the channels integrate across each element in the pattern) and is complete before the action of the normalization network, and thus is not affected by the normalization network. We have confirmed this with numerical calculations (although none are shown here in the interests of space), and it is also evident in the approximate equations approach (see Appendix B). For the normalization model to predict the nonlinear (expansive) tradeoff from the area experiments using grating-patch elements, therefore, the complex channels simply need to have an appropriately expansive function at the intermediate stage in the complex channels. The properties of the normalization network in the model are irrelevant. Similarly, the linear tradeoff for square-element patterns is predicted as a consequence of the linear summation in the simple channels.

Conversely, the predicted compression in the constant-difference experiments at middle to high contrasts (where the normalization effect is strong) is largely determined by the properties of the normalization network itself and does not depend very much on the properties of the complex channels. For example, in the high-contrast predictions of Fig. 14 (right panels), as noted before, the amount of compression shown (the decreased segregatability at the end of the curves) is very much the same for linear complex channels (top right panel) and expansive complex channels (bottom right panel).

In summary, one can find parameter values that allow the normalization model to predict simultaneously: compression in the constant-difference experiments at middle to high contrasts (by choosing the normalization network parameters appropriately) and, for the area experiments in the same contrast ranges, linearity for square-element stimuli (as a result of the linearity of the simple channels) and expansiveness for grating-element stimuli (by choosing the function at the intermediate stage in the complex channels to be appropriately expansive).

3.3.3. Low-contrast expansiveness in constant-difference experiments

We can now return to the predictions of the normalization model for the low-contrast case. As it turns out, for constant-difference experiments done at low contrasts, the function at the complex channels' intermediate stage does intrude to some extent in the normalization model's predictions for stimuli segregated by complex channels. This intrusion is demonstrated in the predictions of Fig. 14 (in the middle column for low contrast, and in the left column for no normalization or alternately, very low contrast). In particular, for complex channels having a piecewise-linear function at their intermediate stage when there is no normalization (top left panel), the ends of each predicted curve are flat — that is, the predicted segregation for all same-sign-of-contrast patterns in a constant-difference series is flat. However, for complex channels with an expansive intermediate stage in the absence of normalization (bottom left panel), there is considerable *increase* in segregation toward the end of the curves. For low amounts of normalization, some difference is still predicted due to the kind of complex channel (the bottom versus top middle panels of Fig. 14). In particular, for expansive but not piecewise-linear complex channels, there is some increase in segregation for same-sign-of-contrast patterns that are just outside the dashed vertical lines at $+45$ and $-45°$ (although at the far ends of the curves there is a downturn in both cases).

In short, at contrasts low enough that the effect of the normalization network itself is still small, the normalization model predicts that the expansiveness in the function at the intermediate stage of the complex channel 'leaks through' and shows up as expansiveness in the constant-difference results.

Thus the normalization model with expansive complex channels not only explains simultaneously the results of area experiments and constant-difference experiments at middle to high contrasts but provides automatically an explanation of a new and initially puzzling result we find in the low-contrast constant-difference experiments here (which we have mentioned previously but postponed talking about until now), namely, the expansiveness seen in some of the results. This expansiveness is particularly clear in the lowest contrast range with the 12 c/deg gratings for observer WS (see lower left panel of Fig. 11). In this panel the curves turn upwards outside the vertical lines at $\pm 45°$ rather than downwards.

This expansive aspect of the results at low contrasts can also be seen in the best-fitting early-local functions for the 12 c/deg grating-element patterns. There it shows up as positive acceleration at low contrasts (slopes greater than 1.0 at low contrasts on the log–log plots of Fig. 13). Notice that for observer WS (upper right panel of Fig. 13), the best-fitting function is expansive at the lowest five or six contrasts; that is, its slope is substantially greater than 1 on the log–log plots. This is also true for the lowest three contrasts for observers JH and CAS. (The fourth observer CV was not tested on contrasts as low as the others.)

While for low-spatial-frequency (3/deg) grating patch elements, there is little if any expansiveness at low contrasts clear to the eye (top left panels top rows Figs. 10 and 11), nonetheless the corresponding best-fitting early-local nonlinearities (not shown here) do show some expansiveness. It is possible that low-spatial-frequency results would also show this upturn clearly even in plots like those of Figs. 10 and 11 if we had studied them at low enough contrasts.

If the expansive function at the complex channels' intermediate stage is the correct explanation of the expansive results at low contrasts in the constant-difference experiments, then there are two other things to say. First, the constant-difference results here suggest that the complex channels' intermediate stage is expansive down to lower contrasts than we knew previously. (The expansiveness in the best-fitting functions in Fig. 13, which indicates expansiveness in the constant-difference results, is visible below the range indicated by the dashed lines, which is the range from the area experiments).

Second, if this explanation is correct, there should *not* be expansiveness at low contrasts in constant-difference results for stimuli segregated by simple channels, e.g. the regularly-spaced square-element patterns used here (although there is some complex-channel intrusion even here as we have discussed in Graham & Sutter, 1996, 1998, and above). This is indeed the case. For example, in the plots of Figs. 8 and 9, the regularly-spaced conditions (top two rows) show very little evidence of expansiveness. Also, the best-fitting functions for regularly spaced squares (e.g. Fig. 12) show little expansiveness at low contrasts. The small amount of expansiveness that is present for regularly-spaced squares may be due to the slight complex-channel intrusion. Consistent with this presumption, there is more evidence of expansiveness in the results for sparsely spaced squares (e.g. bottom row, left panel, Fig. 8) than for regularly-spaced squares, and complex channels may well dominate in the segregation of these sparsely-spaced square-element stimuli (as discussed above).

3.3.4. Two other successes

Finally, we mention two further successes of the normalization model although they concern matters that are side issues here.

The normalization model can naturally predict the kind of difference seen in results at different scales (3 vs. 12 c/deg here, or different sizes of regularly-spaced squares both here and in Graham & Sutter, 1996). The normalization model makes this prediction by scaling the sensitivity to different spatial frequencies in either of two ways: either by assuming a separate sensitivity-setting stage (as was assumed and shown in Fig. 1) or by letting a sensitivity factor be absorbed in each channel.

In addition to the stimuli at the ends of constant-difference series, described here, we have studied another situation for which the normalization model predicts dependence only on contrast ratio and the empirical results show such dependence. This situation is in experiments estimating the bandwidth of simple and complex channels: The two types of elements in a pattern differ in spatial frequency or orientation, and their contrasts are also varied Once the contrasts are high enough, the segregation for a given pattern depends only on the ratio of contrasts in the two element types as is predicted by the normalization model (unpublished results from the study of Graham, Sutter & Venkatesan, 1993).

3.4. Some related issues

We have argued above that the results from constant-difference and area experiments considered together strongly favor the normalization model (Fig. 1) since other known alternatives have been ruled out. In particular, these results considered together are inconsistent even with the relative form of the early-local model (Fig. 2). However, a number of issues might cause some confusion or seem worth further brief consideration.

First, although a relatively-early-local function by itself fails as a model of the intensive nonlinearity, one might wonder if a relatively-early-local function could exist in the full model in addition to the normalization network without undermining the successes of the full

model. The answer is 'yes' but its effects would have to be mild enough in the middle to high contrast range that it did not substantially affect the results of either class of experiment.

Second, it is important in general to distinguish between global contrast-gain controls like that in a normalization network — in which the signal controlling the contrast gain depends on a pooled response representing many orientations and spatial frequencies (but not necessarily a wide spatial area) — and *within-channel* contrast-gain controls which depend only on the response of the channel under consideration and thus only on the spatial-frequency and orientation range to which that channel is sensitive. A within-channel contrast-gain control can, in principle, explain a number of the results frequently attributed to contrast-gain controls, including the saturation seen in response-versus-contrast curves for many physiological and psychological responses, and also the changes in contrast gain resulting from adaptation to patterns. In many situations it would be impossible to distinguish global from within-channel contrast-gain controls. However, although we have not previously addressed this point, a within-channel gain control cannot by itself explain the original phenomenon in our texture-segregation results for which the intensive nonlinearity was invoked. In particular, it cannot explain the compressiveness in the results of constant-difference experiments.

A within-channel gain control (e.g. a late within-channel nonlinear function that acts pointwise on the output of the channels) could exist in our model, however, in addition to a normalization network, because calculations suggest it would have little effect on the results of either kind of experiment described here. There are late within-channel nonlinearities in models of V1 cells by Albrecht and Geisler (1991) and Heeger (see Nestares & Heeger, 1997; Tolhurst & Heeger, 1997a; Tolhurst & Heeger 1997b), where these late nonlinearities are expansive. Also, the pattern-discrimination model of Olzak and Thomas (1998) includes a within-pathway nonlinear function that is both expansive and compressive (in addition to the compression introduced by their divisive global gain control embodied in a normalization network).

A third issue is that of alternate forms of complex channels. We have previously introduced two alternate versions of complex channels that differ from the original complex channels used here in the particular details of the second-stage pooling (Graham & Sutter, 1998). In explaining the area experiments, it did not matter at all which of the three versions of complex channel one considered. However, it does matter to some extent here. One of the versions (version 2 in Graham & Sutter, 1998) does *not* predict expansiveness at low contrasts (like that shown in Fig. 14 bottom row, left and middle panels) and, therefore, is not a satisfactory model for the full set of empirical results. The other two versions make essentially identical predictions for the constant-difference experiments here and so cannot be distinguished by these experiments either.

A fourth and broader issue about complex channels is the question of whether both simple and complex channels are actually necessary. In our previous papers, we have assumed without comment that there were both, as in the diagrams of Figs. 1 and 2. In the motion literature, where both Fourier and non-Fourier channels had similarly been suggested, the existence of two kinds is now a point of contention. Some investigators suggest other models of motion perception altogether, and some suggest that both Fourier and non-Fourier stimuli are processed by a single kind of channel. (Recent reviews of parts of this dispute can be found in Taub, Victor & Conte, 1997 and Clifford & Vaina, 1999). While the answer to this question for static patterns may not be definitive, the evidence suggests that both kinds of channels are probably necessary. One line of evidence is based on luminance and contrast modulation thresholds for modulated white noise (Schofield & Georgeson, 1999). Another is the results of element-arrangement texture segregation experiments measuring the bandwidth of the first stage of complex channels (Graham et al., 1993). These estimated first-stage bandwidths are substantially narrower than those estimated for motion (see Werkhoven, Sperling & Chubb, 1993 and a directly comparable texture experiment reported by Graham, 1994). Thus static texture perception may differ from motion perception in this regard although we suspect that both kinds of channels are necessary in both domains (unless of course one elaborates other stages of the model so much as to effectively embody both kinds of channels in other forms, e.g. Blaloch, Grossberg, Mingolla & Nogueira, 1999).

Finally, the experimental results here show that the normalization network acts quite similarly on the outputs of both simple and complex channels. But there is a related question that we have not yet answered: do complex channels as well as simple channels contribute to the normalization pool? To say it another way, is the total amount of inhibition determined both by complex-channel outputs and simple-channel outputs, or are only the simple-channel outputs determining the inhibition (and therefore the contrast gain)? One recent report (Lu & Sperling, 1996) suggests that the complex (non-Fourier) channels involved in motion perception do *not* contribute to the normalization pool, or, in their terminology, do not contribute to the control of contrast gain. In the experiments reported here with element-arrangement textures, however, the answer to the question of whether complex channels are in the normalization pool is contingent on which version of com-

plex channels (Graham & Sutter, 1998) is correct, and we do not know enough about that yet to have an answer.

4. Summary

The perceived segregation of element-arrangement textures forming *constant-difference-series* of patterns was measured. Textures composed either of grating-patch or of square elements were used. The spatial frequency of the grating patches, and the size and spacing of the square elements, were varied, and a wide range of contrasts was used. These *constant-difference experiments* included conditions and observers that could be directly compared to the previously-published *area experiments* (Graham & Sutter, 1998).

An unexpected result of varying the spacing of square elements in these textures occurred: although patterns composed of regularly-spaced square elements were segregated by simple (linear, Fourier) channels as expected, textures composed of sparsely-spaced square elements were segregated by complex (second-order, non-Fourier) channels. In other words, as small elements get further apart, complex channels generally become better than simple channels at linking those elements together into larger patterns.

The relatively-early-local model of the intensive nonlinearity (Fig. 2) proved untenable since it could not explain simultaneously: (a) the results of the constant-difference experiments reported here — particularly the compressiveness at middle to high contrasts; and (b) the results of previously-published area experiments using the same stimuli in the same contrast ranges: in the area experiments expansiveness is shown for complex-channel stimuli and linearity for simple-channel stimuli.

A contrast-gain control based on inter-channel inhibition — as in a normalization network — can explain simultaneously (a) and (b) as well as a number of other results. A late, within-channel contrast-gain control is not satisfactory as it cannot explain (a). At low enough contrasts of stimuli processed by the complex channels, the normalization model predicts expansiveness rather than compressiveness in constant-difference results. This prediction was confirmed. Thus contrast-gain control as embodied in a normalization network is a tenable and attractive model for the intensive nonlinearity.

Acknowledgements

This research was supported by National Eye Institute Grant EY08459. We are grateful to two anonymous referees and Sabina Wolfson for spending substantial time improving this paper.

Appendix A. Model Equations

A.1. Equations: simple and complex channels

In this subsection, we briefly give the equations for simple and complex channels responses to square-element and grating-element patterns in the absence of any intensive nonlinearity. For further description of why simple and complex channels respond as they do, see Sutter et al. (1989), Graham (1991), Graham et al. (1992), Graham and Sutter (1998).

Let D_S denote the contribution of the simple channels to the predicted perceived segregation for square-element patterns. This contribution is the difference between the relevant simple channels' responses in the checkerboard and striped regions. For the experiments here, great simplicity is introduced because the orientations of the fundamental frequencies in the two regions are so far apart (oblique in checkerboard region, horizontal in striped region). Thus, to a very good approximation, the channels segregating these regions have a (non-zero) response in at most one of the two regions. Further, all the relevant simple channels can be approximated by one channel, which we will call the tuned simple channel, which is characterized by a receptive field matched to the spacing of the pattern elements. (See a comparison of filtered responses from the full model and the equation below in Graham, 1991, and Graham et al., 1992.) Also, the tuned simple channel response R_S pooled across the one region in which it has a non-zero response is approximately proportional to its maximum response in that region, and this is the response from the receptive field centered on a strip of the more effective of the two element types. Since the excitatory center is stimulated by a strip of the more effective elements and the inhibitory surround by a strip of the other elements, the response of the receptive field will just be the difference between its stimulation by the two types of elements. In symbols, therefore,

$$D_S = R_S = w_S \cdot \left| A_1 \cdot C_1 - A_2 \cdot C_2 \right| \quad (1)$$

where the quantities C_1 and C_2 are the (signed) contrasts of the two element types and A_1 and A_2 are the areas of the two element types. The parameter w_S represents the sensitivity of the observer's simple channels to the type of pattern under consideration (i.e. whether it is composed of square-element or grating-elements, the size or spatial frequency of the elements, the orientation of the elements) but w_S does not depend on the contrast or areas of the elements which are represented by the other parameters. For grating-element patterns, we will let $w_S = 0$. (The contribution from the simple channels is so small as to be negligible but is not in fact quite equal to 0).

Let D_X be the contribution of the complex channels to segregation, which, by analogous reasoning to that

described above, can be approximated by the response of the tuned complex channel, which is the difference between its response to the two types of elements. This tuned channel is a complex channel having its first filter sensitive to the spatial frequency and orientation of the grating patches and its second filter sensitive to the fundamental frequency and orientation in either the checkerboard or the striped region. Then:

$$D_X = R_X = w_X \cdot \{A_1 \cdot |C_1|^{k_m} - A_2 \cdot |C_2|^{k_m}\} \quad (2)$$

where C_1, C_2, A_1 and A_2 are as before. The parameter w_X reflects the sensitivity of the complex channels. The exponent k_m describes the expansiveness or compressiveness of the pointwise function at the intermediate stage in the complex channels. (Of course, if the intermediate function were something more complicated than a power function, we would need to replace Eq. (2) by a more complicated expression. Fortunately, for this paper a single power law is sufficiently general). The value of w_X would generally be high for grating-element patterns. But w_X is not necessarily zero for square-element patterns because some complex channels can signal the difference between the checkered and striped regions of square elements and thus can contribute to segregation. (These contributing complex channels are those having a first stage sensitive to high spatial frequencies present in individual square elements — e.g. at the edges — and a second stage sensitive to the fundamental frequency and orientation of the checkerboard or striped arrangements.) In Graham and Sutter (1998) two alternate versions of complex channels were also suggested, which differed from those in Eq. (2) in the details of second-stage pooling. These are mentioned briefly in the discussion section but won't be considered in detail here.

A.2. Equations: combination and decision rules to predict response of observer

Next, the contributions of the complex and simple channels to segregation need to be combined to predict the response of the observer. The full family of combination and decision rules we consider reduces here to the combination:

$$D = \{D_S^{k_d} + D_X^{k_d}\}^{1/k_d} \quad (3)$$

where the exponent k_d is the exponent characterizing the Minkowski pooling that occurs at the Comparison and Decision Rule (near right of Figs. 1 and 2). The observer's rating of perceived segregation is assumed to be a monotonic function of this predicted value D. (See illustration in Fig. 16 of Graham et al., 1992).

Note that many different perceptual processes are presumably represented by this very simple decision stage as there are undoubtedly many processes beyond the channels and intensive nonlinear processes explicitly considered in our models (e.g. He & Nakayama, 1994; Lennie, 1998). Indeed, a number of higher processes that may act on element-arrangement textures like those used in the studies here have recently been suggested (Beck, 1994; Pessoa, Beck & Mingolla, 1996; Grossberg & Pessoa, 1998). For our purposes, however, elaborating these higher-level stages in anything beyond this simple decision rule has not proved necessary. For more discussion of this stage, see our earlier papers (Sutter et al., 1989; Graham et al., 1992; Graham & Sutter, 1998, Fig. 4 and Appendix).

A.3. Equations: early, local models

Intuition into why an early, local nonlinear function leads to the correct predictions for the same-sign-of-contrast patterns can be found in Graham (1991, Fig. 18.8 and 18.9) and in Graham et al. (1992, Figs. 15 and 16). A range of predictions is shown in Graham and Sutter (1996, Fig. 6).

To calculate approximate predictions from the original strong form of an early-local model (assuming that the early, local nonlinear function acted directly on the stimulus) one can use Eqs. (1) and (2) above, but substitute the outputs of the early local pointwise nonlinear function, which will be called $r(C_i)$, in place of the original contrasts in those equations. The resulting quantities will be subscripted ELN to distinguish them from the corresponding quantities in equations above:

$$D_{S,\,ELN} = R_{S,\,ELN} = w_S \cdot |A_1 \cdot r(C_1) - A_2 \cdot r(C_2)| \quad (4)$$

$$D_{X,\,ELN} = R_{X,\,ELN} = w_X \cdot |A_1 \cdot |r(C_1)|^{k_m} - A_2 \cdot |r(C_2)|^{k_m}| \quad (5)$$

To predict the observer's response (up to a monotonic transformation) combine the two overall differences and call the result D_{ELN}:

$$D_{ELN} = \{D_{S,\,ELN}^{k_d} + D_{X,\,ELN}^{k_d}\}^{1/k_d} \quad (6)$$

To compute predictions from the *modified* form of the early-local model, that is, the relatively-early-local model of Fig. 2, one in principle needs to filter the stimulus with the sensitivity-setting stage's filter (the filter near the left of Fig. 2) and then proceed as with the original model. Approximately speaking, however, in keeping with our general approach, one can replace C_i in Eqs. (4) and (5) by $S_i \cdot C_i$ where S_i is the sensitivity of the early sensitivity-setting filter. Thus, using RELN now as the subscript gives:

$$D_{S,\,RELN} = R_{S,\,RELN} = w_S \cdot |A_1 \cdot r(S_1 \cdot C_1) - A_2 \cdot r(S_2 \cdot C_2)| \quad (7)$$

$$D_{X,\,RELN} = R_{X,\,RELN} = w_X \cdot |A_1 \cdot |r(S_1 \cdot C_1)|^{k_m} - A_2 \cdot |r(S_2 \cdot C_2)|^{k_m}| \quad (8)$$

A.4. Equations: normalization (inhibition among channels)

For the normalization model we need to consider the 'other' channels that contribute to the normalization pool (because they respond well in both the checkerboard and the striped region) although they do not contribute to segregation itself (because they respond equally well in the two regions). For the square elements, many of these 'other' channels will be simple channels responding to the high spatial frequencies in the edges of the squares; thus their responses are dependent more on the perimeter length rather than on the area of the elements. For the grating elements, many of the 'other' channels will be simple channels responding to the high spatial frequencies of the grating patches.

The spatially-pooled regional responses of these 'other' simple channels is approximately equal to (Graham et al., 1992):

$$R_{OS} = \{|w_{O1} \cdot C_1|^{k_{sp}} + |w_{O2} \cdot C_2|^{k_{sp}}\}^{1/k_{sp}} \quad (9)$$

where k_{sp} is an exponent describing pooling across spatial position within the output from any single channel, and w_{O1}, w_{O2}, are parameters describing the sensitivity of the 'other' simple channels to elements of type 1 or type 2. This equation can be simplified for the constant-difference experiments, since both elements have identical spatial characteristics and thus $w_{O1} = w_{O2} = w_{OS}$, to

$$R_{OS} = w_{OS}\{|C_1|^{k_{sp}} + |C_2|^{k_s}\}^{1/k_{sp}} \quad (10)$$

It is also possible that some 'other' channels are complex channels. One possible set of complex 'other' channels would be those having first filters sensitive to the square-element edges or to the grating elements' spatial frequency and having second filters sensitive to the frequency of alternation between elements and inter-element spaces; this frequency is twice the frequency of alternation between the two types of elements. If complex channels contribute to the normalization pool (a topic we consider in the discussion), the possibility of complex 'other' channels can be represented as follows where w_{OX} is the sensitivity of the 'other' complex channels to the elements of the pattern under consideration.

$$R_{OX} = w_{OX} \cdot \{|C_1|^{k_{sp}k_m} + |C_2|^{k_{sp}k_m}\}^{1/k_{sp}} \quad (11)$$

The normalization pool will consist of R_X, R_S, R_{OS}, and R_{OX}. These channel responses are combined by Minkowski pooling with the exponent k_n. (The predictions are not particularly sensitive to this exponent, so we generally use $k_n = 2$ which makes our normalization pool much like that in the models of Heeger, 1991, 1992a,b and others.) Letting POOL denote the normalization pool,

$$\text{POOL} = \{\sigma + R_S^{k_n} + R_X^{k_n} + R_{OS}^{k_n} + R_{OX}^{k_n}\}^{1/k_n} \quad (12)$$

where a constant σ provides a necessary non-zero minimum for the expression and sets the extent of the linear range. Any imbalance among the four types of channels in contributing to the pool can be easily incorporated. As explained just before Eq. (1), $D_S = R_S$ for the experiments here, and similarly $D_X = R_X$. Thus POOL can be written equivalently as:

$$\text{POOL} = \{\sigma + D_S^{k_n} + D_X^{k_n} + R_{OS}^{k_n} + R_{OX}^{k_n}\}^{1/k_n} \quad (13)$$

Then $R_{S,\,NORM}$, the response of a simple channel after normalization has been applied at the spatial position yielding the peak response, will simply be the response before normalization at that same position divided by the normalization pool's response at that position:

$$R_{S,\,NORM} = R_S/\text{POOL} \quad (14)$$

and similarly for complex channels:

$$R_{X,\,NORM} = R_X/\text{POOL} \quad (15)$$

To finish the prediction from the normalization network,

$$D_{NORM} = \{D_{S,\,NORM}^{k_d} + D_{X,\,NORM}^{k_d}\}^{1/k_d} \quad (16)$$

which can be written equivalently in a number of ways, one of which is:

$$D_{NORM} = \frac{\{D_S^{k_d} + D_X^{k_d}\}^{1/k_d}}{\{\sigma + D_S^{k_n} + D_X^{k_n} + R_{OS}^{k_n} + R_{OX}^{k_n}\}^{1/k_n}} \quad (17)$$

In our previous presentations of this equation we assumed that the pooling exponents for normalization and decision (k_n and k_d) were identical. This need not be true, and it is perhaps clearer to preserve the distinction as we have done in Eq. (17) although this distinction is not one that our results to date are sensitive enough to discriminate. Previous publications also assumed only a single kind of 'other' channels since we had not yet confronted the fact of expansiveness at the intermediate stage of complex channels.

Concealed in the derivation as presented above is the question of how responses from different spatial positions within a channel are combined both in the normalization pool and at the decision stage (that is, the action of the combination rules across space at both these stages). To explore this, one can go back and explicitly include spatial pooling and allow exponents describing those two kinds of spatial pooling to vary. As it turns out, the decision-stage and normalization-stage spatial-pooling exponents get buried within constants multiplying the numerator and denominator respectively in Eq. (17). Thus they have no effect here (since the observer's rating is assumed to be a monotonic function of the predicted value D_{NORM}). In fact, all the parameter values characterizing the normalization and decision stages seem to be unimportant for our arguments in this paper.

Appendix B. Some explanations of model predictions

B.1. Why should complex channels be more involved in segregating the sparsely-spaced squares than the regular squares?

Consider the *large-regular* and the *small-sparse* patterns. They have the same fundamental frequency (the center-to-center spacing of the elements is the same), so the simple channels that respond best are the same for both patterns. Also at least some of the complex channels that respond significantly will be the same for both the large-regular and small-sparse patterns. (These are complex channels having first-stage filters that are sensitive to very high spatial frequencies — those frequencies represented in the edge of the square elements of whatever size square — and second-stage filters that are sensitive to the fundamental frequency and orientation in either the checkerboard or the striped region). Consider what happens to the responses of these two classes of channel as one shifts from the large-regular to small-sparse patterns (as the square size is decreased while keeping the center-to-center distance between squares the same). The simple-channel response decreases in proportion to the area of the elements (proportional to width-squared, a factor of 16 in the current experiments), but the complex-channel response only decreases in proportion to the total edge length (proportional to width, a factor of 4 in the current experiments). The picture is complicated by the potential presence of a wide variety of complex channels, but it seems clear that the ratio of complex-to-simple-channel response will generally increase substantially as one goes from regular to sparse spacing.

B.2. Why the normalization model predicts compresssiveness in constant-difference experiments

The normalization model easily explains the decreased segregatability for same-sign-of-contrast patterns (the downturn at the ends of the curve for a constant-difference series, in, e.g. Fig. 7 bottom row).

It is due to increased inhibition (a larger denominator in Eq. (17)) in the presence of approximately constant excitation (a constant numerator) as one moves toward the ends of the constant-difference series.

For square element patterns, the effect in Eq. (17) is particularly easy to see. Remember that the response of the simple channel segregating these patterns — before the normalization network — is the same for all patterns in a constant-difference series (see Eq. (1)) and hence the numerator of Eq. (17) stays constant. It turns out the response of the 'other' channels that contribute to the normalization pool only (primarily simple channels responsive to the high-spatial-frequency edges of individual elements) is much larger at the ends of the constant-difference series than in the middle because the contrasts of both element types become very large at the ends. See the stimulus profiles in Fig. 5. This effect is obvious in Eqs. (9) and (10): the value of R_{OS} and/or R_{OX} becomes very large because both $|C_1|$ and $|C_2|$ become very large. Hence the size of the normalization pool — the denominator of Eq. (17) — becomes very large toward the ends of the constant-difference series.

The argument for the case of grating-element patterns and complex channels is complicated by the potential presence of the expansive or compressive function at the intermediate stage (see Eq. (2)) which means that the responses of the complex channels will not necessarily be constant toward the end of the constant-difference series (although they will for an exponent of one). But numerical calculations show that once the contrast is high enough to bring the normalization into action, the effect of the complex channel intermediate stage (even if it is quite expansive) is overcome in these experiments by the normalization pool's compressive behavior (as discussed in the main text in connection with Fig. 14).

B.3. Why the tradeoff in area experiments is independent of the normalization network

According to the normalization model, the effects in the area experiments turn out to be quite independent of the properties of the normalization network, depending primarily on the properties of the channels. The effects in the area experiments are demonstrated in the minima of curves plotting segregation versus contrast in the elements of one type (one area) when the contrast in elements of the other type (another area) is held constant. (See Graham & Sutter, 1998, for more explanation of how these minima show whether there is expansiveness or compressiveness at the intermediate stage of complex channels.) The tradeoff between area and contrast occurs within individual channels (when the channels integrate across each element in the pattern) and the minima produced by this tradeoff are not affected by the subsequent action of the normalization network. We have confirmed this with numerical calculations, but it is also evident in the approximate equations approach. In terms of the equations, the minimum in a segregation-versus-contrast curve occurs when the areas and contrasts of the individual elements are such that D_S and D_X (see Eqs. (1) and (2)) are 0 and hence the numerator of Eq. (17) is 0. Changing the amount of normalization (i.e. changing the values of σ , R_{OS}, and R_{OX}, which are all positive) only changes the value of the denominator, leaving the numerator 0. Hence changing the amount of normalization will not affect the fact that a particular combination of areas and contrasts produces a minimum. Hence changing the

amount of normalization does not affect the predicted amount of expansiveness or compressiveness demonstrated in area experiments.

Appendix C. Some details about calculating the model predictions and fitting the data

For the model predictions and fits reported in this study, we used the approximate-equations approach we have used before. The equations are repeated above in this appendix.

C.1. Sets of model predictions

The conclusions here are based on calculating predictions from many versions of the relatively-early-local and normalization models and comparing them qualitatively to the experimental results. In these calculations, we varied the properties of the normalization network, and/or the early-local nonlinear function and we also varied the exponent of the nonlinear function at the intermediate stage in the complex channel. We also tried the several versions of second-stage pooling described in Graham and Sutter (1998), but those results are not reported in detail here. In these calculations we generally considered only the complex channel and simple channels that were 'tuned' to the patterns in question, but we ran some simulations to check on contamination from other channels (as we did in Graham & Sutter, 1998). Most of the predictions on which the conclusions are based were done letting the exponents in the normalization and decision stages of the model (both across space and across channels) be equal to 2. A sampling of computations with other exponents at these late stages in the model showed no differences in predictions relevant here. The quantities A_1 and A_2 represent the areas of the two element types. Since they are identical in the constant-difference experiments reported here, we can without loss of generality set them equal to 1.0 and let all the sensitivity differences be absorbed in the parameter w_S. See the final paragraphs of Section 2 in the main text for some further discussion of the effect of the approximations involved in this approach.

C.2. Fitting the models to experimental results

We also did one large set of quantitative fits of models to the results from the constant-difference experiments here. We fit both normalization and relatively-early-local models to all the experimental results both for square elements and grating elements. These fits were done in the same way as the set of fits reported in Graham and Sutter (1996) for previous square-element constant-difference experiments except that, for the grating-element patterns here, the parameter w_S was set equal to zero to save time (since it always came out very near zero when allowed to vary in pilot calculations). For square elements the contributions of simple and complex channels, w_S and w_X, were allowed to vary as before. These fits all assumed a complex channel with a piecewise-linear intermediate stage ($k_m = 1$) and assumed that the sensitivities to the two elements (S_i in Eqs. (7) and (8)) were equal to each other as the elements were identical except in contrast. The generally excellent quality of the fits, and also the patterns of best-fitting parameters (and the interactions among these parameters), were all like those reported in Graham and Sutter (1996) and so will not be described in detail here. However, Figs. 12 and 13 do show the best-fitting early-local functions from these fits as a means of comparing compressiveness. More exactly, these figures plot the best-fitting value of $r(S_1 \cdot C_1)$ in Eqs. (7) or (8) as a function of C_1 (assuming $k_m = 1$).

Using exponents k_m other than 1.0 at the complex-channel intermediate stage does not change the excellent quality of the fits to the constant-difference experiment results for either the normalization model or the relatively-early-local model. Nor does it change the degree of normalization deduced from the fits to that model. But it does change the deduced relatively-early-local nonlinear functions (e.g. Figs. 12 and 13) quite dramatically, as the nonlinear function in the complex channels' intermediate stage acts much like a relatively-early-local nonlinear function so the effects in the experimental results would be the concatenation of the two functions. (Indeed, this is the basis for the rejection of the early-local models as is discussed in great detail in the main text.)

References

Albrecht, D. G., & Geisler, W. S. (1991). Motion selectivity and the contrast-response function of simple cells in the visual cortex. *Visual Neuroscience*, *7*, 531–546.

Beck, J. (1994). Interference in the perceived segregation of equal-luminance element-arrangement texture patterns. *Perception and Psychophysics*, *56*, 424–430.

Beck, J., Graham, N., & Sutter, A. (1991). Lightness differences and the perceived segregation of regions and populations. *Perception and Psychophysics*, 257–269.

Bernardete, E. A., Kaplan, E., & Knight, B. W. (1992). Contrast gain control in the primate retina: P cells are not X-like, some M cells are. *Visual Neuroscience*, *8*, 483–486.

Blaloch, A. A., Grossberg, S., Mingolla, E., & Nogueira, C. A. M. (1999). Neural model of first-order and second-order motion perception and magnocellular dynamics. *Journal of the Optical Society of America A*, *16*, 953–978.

Bonds, A. B. (1989). Role of inhibition in the specification of orientation selectivity of cells in the cat striate cortex. *Visual Neuroscience*, *2*, 41–55.

Bonds, A. B. (1993). The encoding of cortical contrast gain control. In R. M. Shapley, & D. M. Lam, *Contrast sensitivity* (pp. 215–230). Cambridge, MA: MIT Press.

Bowen, R. W., & Wilson, H. R. (1994). A two-process analysis of pattern masking. *Vision Research*, *34*, 645–657.

Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *Journal of Neuroscience*, *17*, 8261–8644.

Clifford, C. W. G., & Vaina, L. M. (1999). A computation model of selective deficits in first and second-order motion processing. *Vision Research*, *39*, 113–130.

Foley, J. M. (1994). Human luminance pattern-vision mechanisms: masking experiments require a new model. *Journal of the Optical Society of America A*, *11*, 1710–1719.

Foley, J. M., & Schwarz, W. (1998). Spatial attention: effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. *Journal of the Optical Society of America A*, *15*, 1036.

Geisler, W. S., & Albrecht, D. G. (1992). Cortical neurons: isolation of contrast gain control. *Vision Research*, *22*, 1409–1410.

Graham, N. (1991). Complex channels, early local nonlinearities, and normalization in perceived texture segregation. In M. S Landy, & J. A. Movshon, *Computational models of visual processing* (pp. 273–290). Cambridge, MA: MIT Press.

Graham, N. (1994). Nonlinearities in texture segregation. In: *Higher-order Processing in the Visual System (Ciba Foundation Symposium No. 184)* (pp. 309–329) Chichester: Wiley.

Graham, N., Beck, J., & Sutter, A. (1992). Nonlinear processes in spatial-frequency channel models of perceived texture segregation. *Vision Research*, *32*, 719–743.

Graham, N., Sutter, A., & Venkatesan, C. (1993). Spatial-frequency- and orientation-selectivity of simple and complex channels in region segregation. *Vision Research*, *33*, 1893–1911.

Graham, N., & Sutter, A. (1996). Effect of spatial scale and background luminance on the spatial and intensive nonlinearities in texture segregation. *Vision Research*, *36*, 1371–1390.

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

Grossberg, S., & Pessoa, L. (1998). Texture segregation, surface representation, and figure-ground separation. *Vision Research*, *38*, 2657–2684.

He, Z., & Nakayama, K (1994). Perceiving textures; beyond filtering. *Vision Research*, *34*, 151–162.

Heeger, D. J. (1991). Computation model of cat striate physiology. In M. S. Landy, & J. A. Movshon, *Computational models of visual processing*. Cambridge, MA: MIT Press.

Heeger, D. J. (1992). Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, *9*, 181–197.

Heeger, D. J. (1992). Half-squaring in responses of cat striate cortex. *Visual Neuroscience*, *9*, 427–433.

Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (1996). Computational models of cortical visual processing. *Proceedings of the National Academy of Science USA*, *93*, 623–627.

Hood, D. A. (1998). Lower-level visual processing and models of light adaptation. *Annual Reviews of Psychology*, *49*, 503–535.

Lennie, P. (1998). Single units and visual cortical organization. *Perception*, *27*, 889–935.

Lu, Z.-L., & Sperling, G. (1996). Contrast gain control in first- and second-order motion perception. *Journal of the Optical Society of America A*, *13*, 2305–2318.

Malik, J., & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, *7*, 923–932.

Nestares, O., & Heeger, D. J. (1997). Modeling the apparent frequency-specific suppression in simple cell responses. *Vision Research*, *37*, 1535–1543.

Olzak, L. A., & Thomas, J. P. (1998). Neural recoding in human pattern vision: model and mechanisms. *Vision Research*, *39*, 231–256.

Pessoa, L., Beck, J., & Mingolla, E. (1996). Perceived texture segregation in chromatic element-arrangement patterns: high-intensity interference. *Vision Research*, *36*, 1745–1760.

Robson, J. G. (1988a). Linear and non-linear operations in the visual system. *Investigative Ophthalmology and Visual Science* (*Supplement*), *29*, 117.

Robson, J. G. (1988b). Linear and non-linear behavior of neurones in the visual cortex of the cat. In *New Insights on Visual Cortex*, *the Sixteenth Symposium of the Center for Visual Science*. University of Rochester, Rochester, NY, June, 1988, Abstract p. 5.

Rohaly, A. M., Ahumada Jr, A. J., & Watson, A. B. (1997). Object detection in natural backgrounds predicted by discrimination performance and models. *Vision Research*, *37*, 3225–3235.

Sengpiel, F., Baddeley, R. J., Freeman, T. C. B., Harrad, R., & Blakemore, C. (1998). Different mechanisms underlie three inhibitory phenomena in cat area 17. *Vision Research*, *38*, 2067–2080.

Shapley, R., & Victor, J. D. (1978). The effect of contrast on the transfer properties of cat retinal ganglion cells. *Journal of Physiology*, *285*, 275–298.

Shapley, R., & Victor, J. D. (1979). Nonlinear spatial summation and the contrast gain control of cat retinal ganglion cells. *Journal of Physiology*, *290*, 141–161.

Shapley, R., & Victor, J. D. (1981). How the contrast gain control modifies the frequency response of cat retinal ganglion cells. *Journal of Physiology*, *318*, 161–179.

Schofield, A. J., & Georgeson, M. A. (1999). Sensitivity to modulations of luminance and contrast in visual white noise: separate mechanisms with similar behavior. *Vision Research*, *39*, 2697–2716.

Simoncelli, E. P., & Heeger, D. J. (1998). A model of neuronal responses in visual area MT. *Vision Research*, *38*, 743–761.

Simoncelli, E. P., & Schwartz, O. (1998). Derivation of a cortical normalization model from the statistics of natural images. *Investigative. Ophthalmology and Visual Science (Supplement)*, *39*, Abstract # 1977.

Smith, D. R. R., & Derrington, A. M. (1996). What is the denominator for contrast normalisation? *Vision Research*, *36*, 3759–3766.

Snowden, R. J., & Hammett, S. T. (1998). The effects of surround contrast on contrast thresholds, perceived contrast, and contrast discrimination. *Vision Research*, *38*, 1935–1945.

Solomon, J. A., Sperling, G., & Chubb, C. (1993). The lateral inhibition of perceived contrast is indifferent to on-center/off-center segregation, but specific to orientation. *Vision Research*, *33*, 2671–2683.

Sperling, G. (1989). Three stages and two systems of visual processing. *Spatial Vision.*, *4*, 183–207.

Sutter, A., Beck, J., & Graham, N. (1989). Contrast and spatial variables in texture segregation: testing a simple spatial-frequency channels model. *Perception and Psychophysics*, *46*, 312–332.

Taub, E., Victor, J. D., & Conte, M. M. (1997). Nonlinear preprocessing in short-range motion. *Vision Research*, *37*, 1459–1470.

Teo, P. C., & Heeger, D. J. (1994). Perceptual image distortion. In: B. E. Rogowitz, & J. P. Alleback, *Human vision*, *visual processing*, *and digital display V*, *Proceedings of SPIE*, *2179*, 127–139.

Tolhurst, D. J., & Heeger, D. J. (1997). Contrast normalization and a linear model for the directional selectivity of simple cells in cat striate cortex. *Visual Neuroscience*, *14*, 19–25.

Tolhurst, D. J., & Heeger, D. J. (1997). Comparison of contrast-normalization and threshold models of the responses of simple cells in cat striate cortex. *Visual Neuroscience*, *14*, 293–309.

Victor, J. D, Conte, M. M., & Purpura, K. P. (1997). Dynamic shifts of the contrast-response function. *Visual Neuroscience*, *14*, 577–587.

Watson, A. B., & Solomon, J. A. (1997). Model of visual contrast gain control and pattern masking. *Journal of the Optical Society of America A*, *14*, 2379–2391.

Werkhoven, P., Sperling, G., & Chubb, C. (1993). The dimensionality of texture-defined motion: a single channel theory. *Vision Research*, *33*, 463–485.

Wilkinson, F., Wilson, H. R., & Ellemberg, D. (1997). Lateral interactions in peripherally viewed texture arrays. *Journal of the Optical Society of America A*, *14*, 207–2067.

Wilson, H. R. (1993). Nonlinear processes in visual pattern discrimination. *Proceedings of the National Academy of Sciences USA*, *90*, 9785–9790.

Wilson, H. R., & Humanski, R. (1993). Spatial frequency adaptation and contrast gain control. *Vision Research*, *33*, 1133–1149.

Zetzsche, C., Krieger, G., Schill, K., & Treutwein, B. (1998). Natural image statistics and cortical gain control. *Investigative Ophthalmology and Visual Science (Supplement)*, *39*, Abstract # 1978.