Physica Scripta. Vol. 39, 147–152, 1989.

Low-Level Visual Processes and Texture Segregation

Norma Graham

Department of Psychology, Columbia University, New York, New York 10027, U.S.A.

Received February 27, 1986; accepted February 1 1988

Abstract

Parallel processing by a set of multiple mechanisms selectively sensitive along various dimensions of visual patterns – e.g., spatial frequency, orientation, direction-of-motion – seems to be one of the earlier states in the perception of form and objects. The evidence for these mechanisms – which comes primarily from neurophysiology and psychophysics using near-threshold patterns (patterns of such low contrast that they are imperfectly discriminable from a blank field of the same mean luminance) – and the properties of these mechanisms are briefly reviewed. The responses of these mechanisms to some texture patterns used to investigate the segmentation of the visual field into separate regions is illustrated.

1. Multiple mechanisms – low-level processing of visual patterns

Parallel processing by a set of multiple mechanisms selectively sensitive along a number of dimensions of visual patterns (e.g., spatial frequency, orientation) is thought by many to comprise one of the earlier states in the perception of forms and objects. These mechanisms are thought to be an early stage in pattern vision, but not the earliest. They are probably not in the retina or lateral geniculate nucleus, for example, because they exhibit orientation-selectivity. On the other hand, they are certainly not very high level (e.g., not at the level of object perception) for their behavior depends on proximal (retinal) stimulus properties rather than distal properties (e.g. they do not exhibit size constancy). The physiological substrate for these mechanisms might conceivably be the neurones of cortical areas V1 and V2; see, for example, Refs. [1–4].

Many words in addition to "mechanism" have been used in the visual literature for this same concept, e.g., detector, analyzer, channel, pathway. I will use "mechanism" here to mean something analogous to a single neuron in that its output is a single number (which might be instantaneous firing rate) at each moment in time. The word "channel" will be reserved here for a collection of such mechanisms that are homogeneous in a sense to be described later.

A good deal of the best evidence for these mechanisms has been inferred from human observer's detection and identification of near-threshold spatiotemporal patterns. (Near-threshold patterns are patterns of such low contrast that they are imperfectly discriminable from a blank field of the same mean luminance.) Introduction to and a review of this evidence has already been published [5–6] and only a summary will be repeated here. (Reviews of multiple mechanisms including evidence from suprathreshold pattern vision can be found, for example, in Refs. [6, 9–11].

1.1. *Simplifying assumptions*

How does a physchophysicist study a piece out of the middle of a visual system? As is true in all science, whether one makes it explicit or not, one makes simplifying assumptions in order to isolate one part from the rest. And one justifies these assumptions by their success in explaining phenomena, by their contributions to our understanding. One has to be judicious, of course, in making assumptions and one has to constantly worry about their implications. One has to figure out ways of testing different assumptions independently as much as possible. But one is always making assumptions.

The assumptions made by a number of psychophysicists to study these mechanisms are of two kinds.

Simplifying assumption 1: Bypass early levels of the visual system by specifying a simple transformation (usually linear) between the stimulus on the retina and the mechanism output. Keeping the mean (space- and time-average) luminance of the patterns constant and keeping the contrast of the patterns low makes linearity more plausible by reducing the effects of, for example, retinal light adaptation processes which are highly nonlinear (see, for, example Ref. [12]).

Simplifying assumption 2. Incorporate all higher levels of the visual system into a simple decision rule: e.g., an observer says he sees a pattern (rather than a blank field) if and only if output of maximally-responding mechanism is above some criterion. In the near-threshold detection and identification tasks, such simple decision rules may work as well as they do because these are tasks in which an observer can be "correct" and can, therefore, be trying to maximize his proportion correct. Optimal behavior is well described by such decision variables in at least some cases (see, for example, Ref. [13]).

1.2. *Conclusions about multiple mechanisms*

Interpreted with care, the evidence from near-threshold psychophysical experiments (buttressed by physiological evidence) suggests that multiple mechanisms, selectively sensitive to different ranges of value, exist along at least the following dimensions: spatial frequency, orientation, spatial position, and direction of motion at velocities higher than 1 or 2 deg/s. These multiple mechanisms also seem to have "labelled outputs", that is, higher visual processes can tell which mechanism a given output came from. Analogous experimental results on the temporal position dimension can be explained by assuming that each mechanism's output at a particular time depends only on the recent past and is labelled. For the temporal-frequency dimension (at any fixed spatial frequency), evidence suggests that there are not narrowly tuned mechanisms although there may be some broadly tuned ones and they may have labelled outputs. For the eye-of-origin dimension, although the evidence suggests selectively-sensitive mechanisms (at least at some spatial and temporal frequencies), these mechanisms seem not to have labelled outputs.

The next section (and the appendix) gives further details of the spatial properties of these multiple mechanisms.

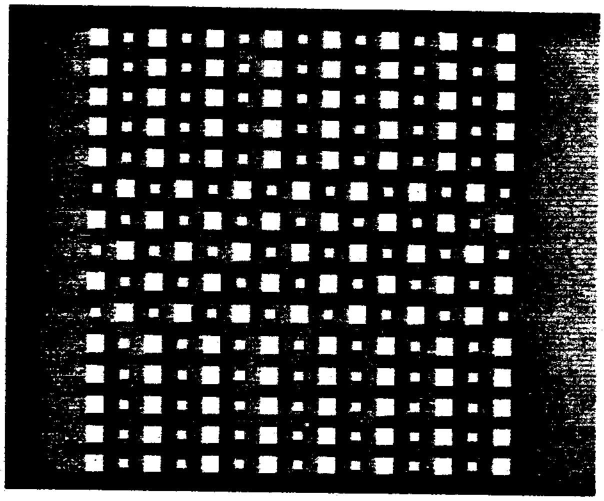

Fig. 1. A pattern in which the two elements have differenent sizes. In the top and bottom regions of the pattern these elements are arranged in stripes; in the middle region they alternate within each row and column giving a checkerboard-type appearance. To get the greatest perceived segregation, you should probably hold this figure so that your thumb at arm's length covers 1/4 to 1/2 of the texture pattern.

Our knowledge of these low-level multiple mechanisms is good enough that it may be helpful to explicitly compute their outputs in response to more complex, suprathreshold patterns – pattern of the kind used by investigators interested in higher-level visual processes. Such computations may help us understand which aspects of the observer's behavior are primarily due to higher-level visual processes and which are heavily determined by the properties of the low-level multiple mechanisms. (To ascribe a locus to some behavior when the whole trivial system is undoubtedly always involved is tricky and can easily lead one into misleading verbalizations. For, even if the output of an early level seems to correspond exactly to the observer's behavior, the higher levels must be doing something. And if you explain the observer's behavior on the basis of the early level processes, you must be assuming – at least implicitly – that the higher levels are not undoing the correspondence between the observer's behavior and the outputs of the early level. For an explicit discussion of this issue see Ref. [14].) The rest of this paper gives an example of the outputs of these multiple mechanisms in response to one kind of suprathreshold pattern. These are from computations that Jake Beck, Anne Sutter, and I have been making in an attempt to further understand what has come to be called "texture segregation" or "texture segmentation" or "immediate and effortless texture discrimination".

2. Texture segregation

When viewing patterns like those in Figs. 1 and 2, observers typically report that they see three regions – a checkerboard-like region in the middle with striped regions at the top and the bottom. (Results from experiments with these patterns are given in Refs. [15, 16].) These are multi-element patterns in which differently appearing regions are made by intermixing appropriately two kinds of elements (large and small squares in Fig. 1, bright and dim squares in Fig. 2). Although multi-element regions as in these patterns are not what are most typically called "textures" in English, such multi-element patterns are referred to as "textures" in the visual literature. The processes leading an observer to see separate regions in such patterns are sometimes thought to be related to what Gestalt psychologists called "grouping" or "figure-ground segregation" processes.

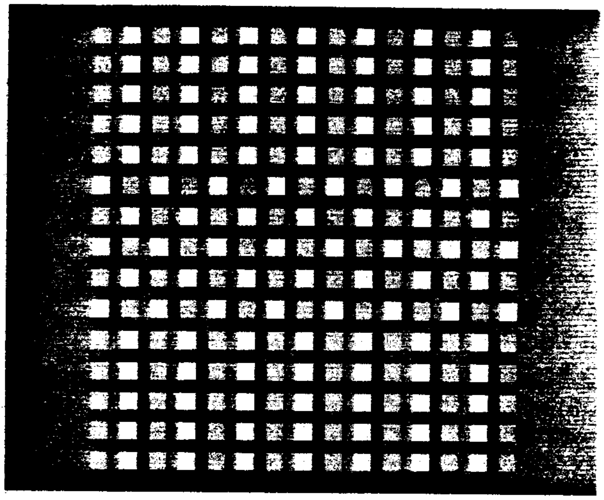

Fig. 2. A pattern like Fig. 1 except the two elements are of the same size but different contrasts.

Texture-segregation experiments are a subject of considerable interest these days, and a number of investigators have discussed both low-level and higher-level visual processes that might be responsible for this kind of segregation (see Refs. [17–33].) And texture segregation does seem to be a promising direction in which to extend the rigorous and quantitative multiple-mechanism models (that do so well for near-threshold psychophysics) to include at least some interesting aspect of higher-level processes. It was clear from early texture-segregation results that the properties to which the low-level multiple mechanisms were thought to be sensitive (e.g., orientation and spatial frequency) seemed to play a major role in determining the extent of texture segregation [34–41]. Further, in the typical texture-segregation experiment, the observer is not supposed to concentrate or "scrutinize" but simply to report a first or immediate impression. Such a task – while requiring some knowledge of higher-level visual stages to explain – might be more tractable than other suprathreshold pattern discrimination tasks where the observer, encouraged to use all the information available, may produce very complicated behavior.

2.1. *Effects of contrast and size of element*

When the difference between the contrasts (Fig. 2) or sizes (Fig. 1) of the squares is decresed in the texture patterns here the perceived segregation reported by observers decreases [17]. Such a result seems reasonable on many grounds and similar results have been reported with other kinds of texture patterns [26, 42].

Less obvious, perhaps, is what should happen when the two kinds of squares differ not only in size but also in contrast. On some points of view (e.g., those that consider size and contrast as two distinct features of visual objects) differentiating the two elements on both size and contrast might be expected to enhance the perceived segregation. After all, the elements now differ from each other in two ways rather than just one.

In fact, however, what happens for any particular pair of sizes depends on contrast and depends nonmonitonically. As the contrast of the smaller squares increases starting at a value equalling that of the large squares, the perceived segregation first decreases (the checkerboard and striped regions seem less and less distinct from each other). After sufficient increases in contrast, however, the perceived segregation starts to increase again [16]. Recent experiments of Anne Sutter's and Jake Beck's have shown that the perceived segregation reaches its minimum when the ratio of the small-centre contrast to the large-square contrast is about the same as the ratio of the large-square area to be small-square area, that is, when the product of contrast times area is about the same for two kinds of squares. (This is true for all the absolute sizes studied although the effect of contrast is more dramatic for some sizes than for others.) It would have been nice to have shown a figure here of an unequal-size unequal-contrast pattern in which the regions did not segregate well perceptually, but contrast in photographs are almost always changed dramatically by publication and – as said above – the extent of perceived segregation depends critically on contrast ratio.

2.2. *Outputs of low-level mechanisms sensitive to different spatial frequencies and orientations*

The existence of a trade-off between contrast and area seems entirely plausible if one considers the outputs of the low-level visual mechanisms described in the introduction, mechanisms that are selectively-sensitive to spatial frequency, orientation, and position. (We will ignore time, color, depth for the rest of the discussion). These mechanisms are thought to be, approximately, linear systems characterized by weighting functions that look like the receptive fields of neurons in cortical area V1. That is, they have elongated exitatory and adjacent elongated inhibitory regions. The widths of the regions determine the spatial frequencies to which the mechanism responds best, the orientation of the regions determines the orientation to which the mechanism responds best, and, of course, the position of the weighting function in the visual field determines the position to which the mechanism responds best. In what follows, we will use the word "channel" to refer to an array of mechanisms having weighting functions (receptive fields) that are identical in size and shape (in spatial frequency and orientation) but are located at many different positions in the visual field.

We calculated the outputs of channels of this type to patterns like those in Figs 1 and 2. Details of the calculations are given in the appendix. The outputs to the patterns in Figs. 1 and 2 are shown in summary form in Figs 3 and 4 respectively, and the outputs to two more patterns (described below) are shown in Figs 5 and 6. In each of these figures, outputs are shown for channels sensitive to thirteen different spatial frequencies going from 2 to 128 cycles/screen by the square-root of two (from left to right in the thirteen columns) and to three different orientations – vertical, 45 degrees clockwise from vertical, and horizontal (from top to bottom in three pairs of two rows each). Consider, for example, the upper right square in Fig. 3. Brightness at any position in that square represents the output of a mechanism that has a weighting function (a receptive field) sensitive to 128 cycles/screen and a vertical orientation; further, that mechanism's weighting function is centered at a particular position in the visual field. The positions represented in each square in Figs. 3–6 correspond to the positions within one period of a pattern like that in Fig. 1. Squares in the top, third, and fifth rows of Figs. 3–6 show the centermost period from the checkerboard region of the pattern, and squares in the second, fourth, and bottom rows of Figs. 3–6 show the centermost period from the striped region of the pattern.

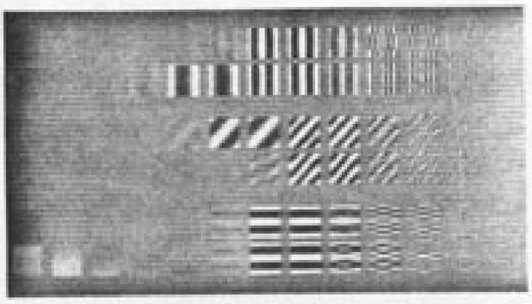

Fig. 3. The outputs of many channels to the texture pattern of Fig. 1. Each square in this figure shows the output of one channel to one period of the texture pattern. The top, third, and fifth rows in this figure show outputs to the middle period of the texture pattern's checkerboard region; the second, fourth, and sixth rows in this figure show outputs to the middle period of the texture pattern's striped region. This figure shows outputs from channels sensitive to thirteen different spatial frequencies (2 to 128 cycles per screen by a factor of square-root of two – in the thirteen columns from left to right) and three different orientations (vertical in top pairs of rows, 45 degrees clockwise from vertical in second pairs of rows, and horizontal in bottom pair of rows).

Different-sized, same-contrast elements. The channel outputs in Fig. 3 are in response to the pattern in Fig. 1 (different-sized same-contrast squares). The low-spatial-frequency channels (left side Fig. 3) show little activity in either the checkerboard or striped region as their receptive fields are so large that they average together several contiguous rows and/or columns of elements. (What activity there is are responses to the edge between the texture pattern and the blank background.) The highest-spatial-frequency channels, on the other hand (right side Fig. 3), show activity near the edges of all the individual elements in the pattern, and the amount of activity is much the same in both the checkerboard and the striped region although distributed differently spatially.

It is only in the range of spatial frequencies corresponding to the periodicity inherent in the pattern that channels show dramatic differences between their outputs in the checkerboard and striped regions. Consider the outputs of the vertical 8 and 11.3 cycle/screen channels (the top two rows, the 5th and 6th column from the left in Fig. 3). The receptive fields in the channel are vertically oriented and their sizes are such that if the excitatory region of a receptive field is centered along one column of squares in the pattern, the inhibitory regions are approximately centered on the neighboring columns of squares. In the striped region, therefore, the receptive fields centered on the big-square columns are getting more excitation (from big squares) and less inhibition (from small squares) than are the receptive fields centered on the small-square columns (which are getting excited by small squares and inhibited by large squares). Hence, in the striped region, this channel shows strongly-modulated activity with high outputs at the center of the big squares in the pattern and low outputs at the center of the small squares. (See second row, fifth and sixth column from left, Fig. 3). In the checkerboard region, however, the vertically-oriented regions of each receptive field are average over different sizes of element (since

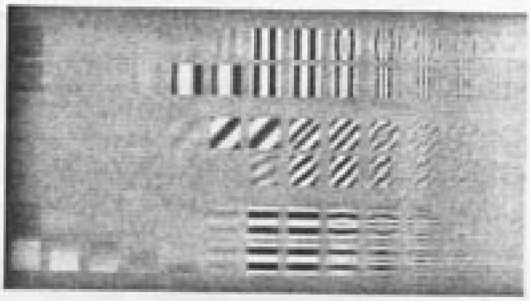

Fig. 4. The outputs of many channels to the pattern of Fig. 2. Other conventions as in Fig. 3.

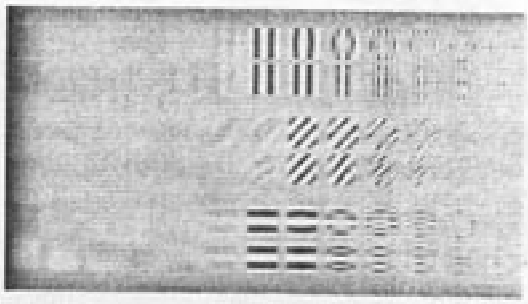

Fig. 6. The outputs of many channels to a pattern like that of Fig. 1 and Fig. 2 except that the two elements differ in both size and contrast in such a way that the product of area times contrast is the same for both.

alternate rows in each column contain big and little squares); and thus there is little modulated activity – the output at each position is about the same.

If one considers the oblique channels (middle two rows Fig. 3) and goes rightward one column (to 11.3 and 16 c/screen), one now sees heavily modulated outputs in the checkerboard region (third row, six and seventh column from the left) and little response in the striped region (fourth row, same column).

At slightly higher spatial frequencies than we have just been discussing, all the channels outputs looks like striped patterns. (Look at vertical and horizontal 16 and 22.6 c/screen channels – the seventh and sixth columns from the right; and at oblique 22.6 and 32 c/screen – the sixth and fifth columns from the right.) This occurs because the receptive fields are now somewhat smaller, and, when the excitatory region is centered over a column (or row, or diagonal) of squares the inhibitory region is seeing almost entirely the spaces between squares.

Same-size, different-contrast elements. The channel outputs in Fig. 4 are in response to the pattern of Fig. 2 (same-size different-contrast squares). They look very much like those in Fig. 3 and for very much the same reasons. Most importantly, in Fig. 4, the vertically-oriented 8 and 11.3 c/screen channel and the obliquely-oriented 11.3 and 16 c/screen channels have heavily modulated responses in one region but not in the other due to different brightnesses of (same-sized) squares whereas in Fig. 3 it was due to different sizes of (same-brightness) squares.

Same-sized same-contrast elements. Figures 5 and 6 show the channel outputs in response to two more patterns. Figure 5 is a control, in effect, for it shows the outputs to a pattern in which all the squares of of the same size and of the same contrast; thus there is actually no difference between the checked and the striped regions that the observer could possibly use to segregate them. Each channel's output is the same in the checked and in the striped region, therefore, (except for effects of the lower edge visible in the responses of the horizontal channels at low spatial frequencies). Notice particularly that the 8 and 11.3 c/screen vertical and the 11.3 16 c/screen oblique channels – the channels matched to the periodicity of the pattern – show no modulated activity in response to this control pattern because all the receptive fields, no matter which column (or row, or diagonal) they are centered on, are getting almost exactly the same stimulation since all the squares are the same.

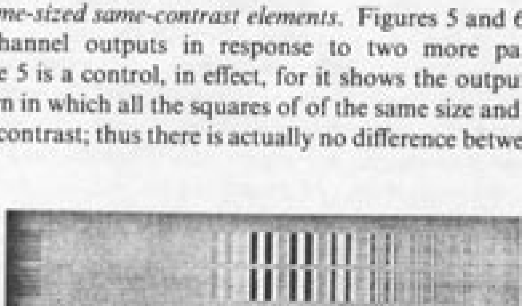

Fig. 5. The outputs of many channels to a control pattern like those of Fig. 1 and Fig. 2 except that both elements are the same size and the same contrast. Other conventions as in Fig. 3.

Different-sized different-contrast elements. Figure 6 is the interesting figure. It shows the channel outputs in response to a pattern in which the squares differ both in size and in contrast but the products of area and contrast are equal – in particular, the large squares have four times the area of the small squares, and the small squares have four times the contrast of the large squares. Look now at the 8 and 11.3 c/screen vertical and the 11.3 and 16 c/screen oblique channels – the ones matched to the periodicity of the pattern. As in the previous figure for the control pattern, they show no modulated response to either the checkerboard or the striped region at all. For the greater luminance has balanced out the smaller size in the integration performed by these mechanisms. (That the squares were really different sizes can be verified by the differences between the checked and striped regions in the patterning of outputs in the high-spatial-frequency channels.)

2.3. *Discussion*

Thus, with equal-contrast unequal-size (Fig. 1) and equal-sized unequal-contrast (Fig. 2) texture patterns, the observer immediately perceives separate regions and some channels produce outputs that are heavily modulated in one region but not in the other (Figs. 3 and 4). with the equal-size equal-contrast control and, more importantly, with the unequal-size unequal-luminance (but equal size × contrast) texture pattern, the observer does not immediately perceive separate regions and no channel produce outputs that are heavily modulated in one region but not in the other (Figs. 5 and 6).

To turn the outputs of these multiple channels into the full-blown perception of the typical observer obviously requires higher levels of visual processing than explicitly described here. (A complete theory would have to specify, for example, just how the boundary between the two regions is found, and why one region looks so "striped" and the other so "checkerboard-y.") The tradeoff between contrast and size, however, seems likely to be a result of the linear integration properties of low-level multiple mechanisms. When contrast and size balance out in the pattern, there are no channels which can produce heavily modulated responses in one region but not in the other (Fig. 6).

The different distribution of activity across positions in the high-spatial-frequency channel outputs (righthand columns) – and these remain even when contrast and size balance out (Fig. 6) – could be and presumably are used by the observer to make discriminations of other sorts. They may, for example, underlie an observer's ability to tell that two kinds of squares are present in these texture patterns even when the observer does not report the regions as segregating. But these spatial distributions seem to play at most a secondary role in determining whether immediate (effortless) texture segregation occurs or not.

Other results. A number of other experimental results using the kind of patterns in Figs. 1 and 2 may well be easily interpretable on the basis of the low-level multiple-mechanism outputs. We are currently considering a number of these. For example, the effect of the absolute size of the squares (recently measured by Anne Sutter and Jacob Beck) seems likely to be due to the greater sensitivity of human observers to certain spatial frequencies than to others (a fact ignored for simplicity's sake above). The effects of the inter-square spacing relative to square size [16] are another likely candidate. To explain differences on dark and light backgrounds and effects when squares differ greatly from the background in luminance requires dropping the linearity assumption and considering the well-known light-adaptation processes that occur even earlier than the mechanisms considered here [12]. An intriguing result – that may not fit into the framework of low-level mechansism here without requiring separate consideration of "on" and "off" mechanisms – is the dramatic effect of having equal-size squares of different luminances where one luminance is less than and one larger than the background [16].

Tentative conclusion. We intend to explore the correlations between channel outputs and experimental results more fully and quantitatively. But suppose for the moment that we did end up concluding – as suggested by the above – that: The extent to which the two texture regions segregate immediately and effortlessly in this pattern is primarily a function of the extent to which some channel or channels produce heavily modulated outputs in response to one region but not to the other.

This tentative conclusion above is not meant to deny the importance of higher-level processes in arranging the final perception of the texture. It is only a description of the information necessary in the stimulus for texture segmentation to occur – a description phrased in terms, however, of the outputs of the lower-level mechanisms well supported by both physiological and psychophysical evidence. That is, it is a simple decision rule – like that used in the very successful models of near-threshold psychophysics (see Section 1.1 above) which bypassses the higher levels.

Nonetheless, I find this tentative conclusion rather disappointing. I would like to find some texture segregation results that could not be so easily understood on the basis of low-level mechanisms, but that were rich and rigorous enough to drive one to specifying in some detail an interesting higher-level "grouping" or "linking" process.

[illegible] that these low-level mechanisms were so narrowband that they respond only to a single spatial frequency and orientation and also continue to ignore (as we have doing) the fact that some mechanisms are more sensitive than others (since observers' sensitivities vary as a function of spatial frequency, orientation, and spatial position). Then the tentative conclusion above could be rephrased as saying: If two texture regions differ sufficiently in the amplitudes of their Fourier Transforms the regions will be seen effortlessly as two separate regions. If, however, they differ only in phase and not in amplitude, the two regions will blend into one [6]. This in turn is a rephrasing of Julesz's original conjecture [38] about the statistical basis of texture discrimination. This conjecture has been challenged on many grounds (see examples in [21, 23]). One wonders how many of these challenges would be overcome if one allows the mechanisms to have some greater-than-zero bandwidth, the kind of bandwidth for example revealed by near-threshold psychophysics and used in the computations here.

In allowing the mechanisms to have some greater-than-zero bandwidth, there is a convergence of the "statistical" approach and the "feature" approach to understanding texture segmentation. This indeed is much the direction Julesz and Beck and their colleagues have taken – see, e.g., Refs. [16, 19, 23, 24, 43].

Turner [46] has, in fact, used Gabor filters to analyze some of the earlier texture patterns used by Julesz and Beck. Two other recent models of texture segregation are Klein and Tyler's [47] autocorrelation approach to texture and related pattern discriminations (including a discussion of the extension of Julesz-type statistics to continuous gray-level patterns), and Grossberg and Mingolla's [48] cooperative-competitive network analysis of texture segregation and other perceptual grouping phenomena.

Appendix. Details of calculations

A convenient (but by no means the only) mathematical expression with which to describe the weighting function $w(x, y)$ used in calculating the outputs is a two-dimensional Gabor function [44, 45]. In one direction it is a Gaussian function multiplied by a sinusoid and in the perpendicular direction (say y) it is simply a Gaussian: We use

$$w(x, y) = m \exp\left(-\ln 2\langle(x - x_0)/W_x\rangle^2 \cos\langle 2\pi f(x - x_0) + \theta\rangle \exp\left(-\ln 2\langle(y - y_0)/W_y\rangle^2\right),$$

where (x_0, y_0) is the center of the weighting function (receptive field), f is its spatial frequency or the reciprocal of the width of an excitatory regionplus an inhibitory region, W_x and W_y give the spatial extent of the field perpendicular and parallel to the regions (equalling the full-bandwidth at half-peak-amplitude in this particular form of Gaussian), θ gives the symmetry of the field (a θ of zero giving even-symmetry), and m gives the height or peak amplitude of the weighting function. Although not explicitly represented in the above equation for simplicity's sake, the orientation can be changed too (by, for example, rotating the x and y coordinates).

The Fourier Transforms of Gabor functions are just Gaussians centered at the appropriate (two-dimensional) spatial frequency – that is, at the spatial frequency and orientation of the weighting function.

The calculations here were done using Gabor functions having parameters modelled after those of Watson (Ref. [45])

although our model is less complete than his in a number of respects.

The spatial extents parallel and perpendicular to the bars – W_x and W_y above – were set equal, thereby producing an orientation half-amplitude full-bandwidth of 38 degrees of rotation.

The spatial extent perpendicular to the bars were chosen so that the spatial-frequency half-amplitude full-bandwidth was one octave ($W_x = W_y = 2/3f$).

The fields had even symmetry ($\theta = 0$).

The heights of the weighting functions in any one channel were equal at all spatial positions (i.e., visual field non-uniformity was ignored).

Also the heights of weighting functions in different channels were chosen so that the peak sensitivities of all the channels were the same (i.e., the non-uniform sensitivity of the observer to different spatial frequencies and – much less dramatically – to different orientations was ignored).

The two-dimensional Fast Fourier Transforms were done over 256 × 256 points. Where we say "*n* cycles per screen" above, we mean *n* cycles per 256 points. One period of the texture patterns was 28 × 28 points, that is, the center-to-center spacing between two rows or two columns of the texture patterns was 14 points.

The displays of Figs. 3–6 were normalized so that – within any one figure – the smallest output produced the lowest and the largest output the highest-possible luminance on the display screen. (These outputs always occur in relatively high spatial frequency channels responding to the edges of the elements and are thus approximately proportional to the highest luminance present in a square.) Inbetween output values were taken linearly into brightness on the display screen but have presumably been distorted by the various reproductive processes intervening between the display screen and the figure here. An output of zero produced a gray in Figs. 3–6 that is the same as the gray in the borders around the squares.

Acknowledgement

This work was partially supported by an NIMH post-doctoral award to Norma Graham and by an Air Force Office of Scientific Research Grant AFOSR-85-0359 to Jacob Beck.

References

1. DeValois, R. L., Albrecht, D. G. and Thorell, L. G., Vision Research **22**, 545 (1982).
2. DeValois, R. L., Yund, E. W. and Hepler, N., Vision Research **22**, 531 (1982).
3. Foster, K. H., Gaska, J. P., Nagler, M. and Pollen, D. A., Journal of Physiology **365**, 331 (1985).
4. Shapley, R. and Lennie, P., Annual Reviews of Neuroscience **8**, 547 (1985).
5. Graham, N. Spatial-frequency channels in human vision: detecting edges without edge detectors, In Visual Coding and Adaptability (Edited by C. S. Harris), pp. 215–262, Lawrence Erlbaum Associates, Hillsdale, New Jersey (1980).
6. Graham, N., Psychophysics of spatial-frequency channels, In Perceptual Organization (Edited by M. Kubovy and J. Pomerantz), pp. 1–26, Lawrence Erlbaum Associates, Hillsdale, New Jersey (1981a).
7. Graham, N., The visual system does a crude Fourier analysis of patterns, In Mathematical Psychology and psychophysiology (Edited by Grossberg). SIAM-AMS Proceedings, Vol. **13**, pp. 1–16, American Mathematical Society, Providence, Rhode Island (1981b).
8. Graham, N., Journal of the Optical Society of America A **2**, 1468 (1985).
9. Braddick, O., Campbell, F. W. and Atkinson, J., Channels in vision: basic aspects, In Handbook of Sensory Physiology (Edited by R. Held, H. W. Leibowitz and H. L. Teuber), Vol. **VIII**, pp. 3–38 (1978).
10. DeValois, R. L. and DeValois, K. K., Annual Reviews of Psychology **31**, 309–341 (1980).
11. Julesz, B. and Schumer, R., Annual Reviews of Psychology **32**, 575 (1981).
12. Shapley, R. and Enroth-Cugell, C., Visual adaptation and retinal gain controls, In Progress in Retinal Research (Edited by N. N. Osborne and G. J. Chader, Vol. 3, pp. 263–343. Pergamon Press, Oxford (1984).
13. Nolte, L. W. and Jaarsma, D., Journal of the Acoustical Society of America **41**, 497–505 (1967).
14. Teller, D., Locus question in visual science, In Visual Coding and Adaptability (Edited by C. S. Harris), pp. 151–176, Lawrence Erlbaum Associates, Hillsdale, New Jersey (1980).
15. Beck, J., Prazdny, K. and Rosenfeld, A., A Theory of Textural Segmentation, In Human and Machine Vision (Edited by J. Beck, B. Hope and A. Rosenfeld), Academic Press, New York (1983).
16. Beck, J., Sutter, A. and Ivry, R., Computer Vision, Graphics, and Image Processing **37**, 299 (1987).
17. Beck, J., Textural segmentation, In Organization and Representation in Perception (Edited by J. Beck), Lawrence Erlbaum Associates, Hillsdale, New Jersey (1982).
18. Beck, J., Biological Cybernetics **48**, 125 (1983).
19. Caelli, T., Perception and Psychophysics **31**, 149 (1982).
20. Caelli, T., Spatial Vision **1**, 19 (1985).
21. Julesz, B., Spatial nonlinearities in the instantaneous perception of textures with identical power spectra, In The Psychology of Vision (Edited by C. Longeut-Higgins and N. S. Sutherland), Philosophical Transactions Royal Society London, **290**, 83 (1980a).
22. Julesz, B., Spatial frequency channels in one-, two-, and three-dimensional vision: Variations on a theme by Bekesy, In Visual Coding and Adaptability (Edited by C. Harris), Lawrence Erlbaum Associates, Hillsdale, New Jersey (1980b).
23. Julesz, B., Figure and ground perception in briefly presented iso-dipole textures. In Perceptual Organization (Edited by M. Kubovy and J. Pomerantz), Lawrence Erlbaum Associates, Hillsdale, New Jersey (1981).
24. Julesz, B., Trends in Neurosciences **7**, 41 (1984).
25. Kimchi, R. and Palmer, S. E., Journal of Experimental Psychology: Human Perception and Performance **8**, 521 (1982).
26. Nothdurft, H. C., Perception **14**, 527 (1985).
27. Nothdurft, H. C., Perception **14**, 539 (1985).
28. Nothdurft, H. C., Vision Research **25**, 1957 (1985).
29. Pentland, A. P., Artificial Intelligence **28**, 293 (1986).
30. Pentland, A. and Gazis, L., Perception of three-dimensional textures, In Supplement to Investigative Opthalmology and Visual Science, Vol. **25**, No. 3, p. 201 (1984).
31. Sagi, D. and Julesz, B., Science **228**, 1217 (1985).
32. Treisman, A., Computer Vision, Graphics, and Image Processing **31**, 1567 (1985).
33. Zucker, S. W. and Cavanagh, P., Spatial Vision **1**, 131 (1986).
34. Beck, J., Perception and Psychophysics **1**, 300 (1966).
35. Beck, J., American Journal of Psychology **85**, 1 (1972).
36. Harvey, L. O. and Gervais, M. J. Perception and Psychophysics **24**, 534 (1978).
37. Harvey, L. O., Jr. and Gervais, M. J., Journal of Experimental Psychology: Human Perception and Performance **7**, 741 (1981).
38. Julesz, B., IRE Transactions on Information Theory **IT-8**, 84 (1962).
39. Julesz, B., Scientific American **232**, 34 (?).
40. Julesz, B., Perceptual limits of texture discriminination and their implications for figure-ground separation, In Formal Theories of Perception (Edited by E. Leeuwenberg), Wiley, New York (1978).
41. Richards, W. and Polit, A., Kybernetik **16**, 155 (1974).
42. Caelli, T. and Moraglia, G., Vision Research **25**, 671 (1985).
43. Julesz, B. and Caelli, T., Perception **8**, 69 (1979).
44. Daugman, J. G., Journal of the Optical Society of America A **2**, 1160 (1985).
45. Watson, A. B., Detection and recognition of simple spatial forms, In Physical and Biological Processing of Images (Edited by O. J. Braddick and A. C. Sleigh), Springer-Verlag, New York (1983).
46. Turner, M. R., Biologica Cybernetics **55**, 71 (1986).
47. Klein, S. A. and Tyler, C. W., Journal of the Optical Society of America A **3**, 868 (1986).
48. Grossberg, S. and Mingolla, E., Perception and Psychophysics **38**, 141 (1985).