1 Psychophysics of Spatial-Frequency Channels

Norma Graham
Columbia University

ABSTRACT

The concepts of spatial-frequency analysis (Fourier analysis) and of spatial-frequency channels are briefly introduced. The various kinds of psychophysical experiments providing evidence for spatial-frequency channels are discussed at some length. Tentative conclusions from these experiments about properties of the spatial-frequency channels are then described. Finally, several ways in which spatial-frequency analyses might be useful in explaining perceptual phenomena are mentioned.

INTRODUCTION

During the last decade, various investigators have suggested using spatial-frequency analyses of visual stimuli to help understand perceptual phenomena. At the very least, spatial-frequency analyses are enlightening simply as a new description of visual stimuli. In a sense, spatial-frequency descriptions and pointwise descriptions are opposites. A stimulus consisting of a single spatial frequency is completely localized on the frequency dimension but infinitely extended in space, while a stimulus consisting of a single point is completely localized in space but infinitely extended on the frequency dimension. Having two such different descriptions that emphasize different aspects of the stimulus may well suggest ideas to the human investigator who is trying to understand visual perception.

In the long run, however, spatial-frequency analyses will be most useful if they help in describing not just the stimuli themselves but the visual system's responses to the stimuli as well. A large number of psychophysical results are now being explained on the assumption that the visual system contains multiple channels, each sensitive to a different range of spatial frequencies. The ability of this concept of multiple spatial-frequency channels to account for a wide range of psychophysical evidence is quite impressive. Perhaps it will be able to account for perceptual phenomena as well.

The psychophysical evidence for the existence of spatial-frequency channels is less clear-cut, however, than some investigators seem to believe, and the probable properties of the channels are somewhat different from those frequently assumed in discussions of visual perception. The purpose of this chapter is to introduce and discuss the body of psychophysical evidence regarding spatial-frequency channels in the hope that a clearer understanding of this evidence will aid in developing and evaluating explanations of visual perception that depend on spatial-frequency analyses.

First, the concepts of spatial-frequency analysis (that is, Fourier analysis) and of spatial-frequency channels are briefly introduced. More extensive introduction can be found in a number of places (e.g., Cornsweet, 1970; Graham, 1979; Julesz, 1980a; Weisstein & Harris, 1980a). Then the various kinds of psychophysical experiments providing evidence for spatial-frequency channels are discussed at some length. After this section on psychophysical evidence, the probable properties of spatial-frequency channels in the human visual system (or the possible meanings of the phrase *spatial-frequency channels*) are described. Finally, a very brief section mentions several ways in which spatial-frequency analyses might help in explaining perceptual phenomena.

Fourier Analysis

To understand Fourier analysis as applied to visual patterns, begin by considering one simple pattern—a *sinusoidal grating*. A sinusoidal grating is a pattern that looks like a set of blurry, alternating dark and light stripes. In such a pattern, the luminance in the direction perpendicular to the stripes varies sinusoidally, whereas the luminance in the direction parallel to the stripes is constant. As this chapter is not concerned with color, we generally ignore wavelength, making the assumption that the wavelength composition of the light is approximately the same at every point of a pattern. Several parameters are used in describing sinusoidal gratings. The one of most concern here is *spatial frequency*, which is the number of cycles of sinusoid per unit distance or, in other words, the number of dark-bar–light-bar pairs per unit distance. (The usual unit of distance is a degree of visual angle.) Spatial frequency and size of bar are inversely proportional to each other; a grating of high spatial frequency has narrow bars, and a grating of low spatial frequency has wide bars. The *mean luminance* of a grating is the average of the luminances at every point across the whole spatial extent of the grating. The *contrast* of a grating is a measure of the difference between the maximum luminance and minimum luminance, usually taken to be one-half that difference divided by the mean luminance. One of the attractions of sinusoidal gratings is that the mean luminance of a grating can easily be held constant, keeping the observer in a relatively constant state of light adaptation, while the contrast and spatial frequency are varied.

Now consider any visual stimulus, the luminance of which varies along only one dimension (called a one-dimensional visual stimulus). As follows from a theorem due to the 18th-century mathematician Fourier, *any* such stimulus can be constructed by superimposing sinusoidal gratings of different spatial frequencies, contrasts, and phases (positions). Further, there is only one set of sinusoidal gratings that can be superimposed to form a particular stimulus. Thus, any one-dimensional visual stimulus can be described as containing certain spatial frequencies. The amount and phase of each spatial frequency contained in a stimulus (the contrast and phase of the component sinusoidal grating of that spatial frequency) are given by the *Fourier transform* of the stimulus.

Analysis into spatial frequencies is not restricted to one-dimensional visual stimuli. Consider an ordinary black-and-white photograph as an example of a two-dimensional visual stimulus. Any such pattern can be constructed by adding up sinusoidal gratings that differ in orientation as well as in contrast and phase. The amount and phase of each spatial frequency at each orientation is given by a *two-dimensional Fourier transform*. Fourier transforms can be computed in higher dimensions as well.

For lack of space, this chapter concentrates on one-dimensional patterns differing in spatial-frequency content but identical in orientation. Many experiments have used one-dimensional patterns differing in orientation but identical in spatial-frequency content, and some experiments have used two-dimensional patterns varying both in spatial frequency and orientation. The only discussion of these experiments on orientation occurs in the "Sensitivity Characteristics: Bandwidth" section. The third dimension, depth, is ignored entirely.

Spatial-Frequency Channels

Let's start by considering one extreme kind of spatial-frequency-channel model. Consider a set of channels, each of which is sensitive to a very narrow range of spatial frequencies, a range so narrow that the channel effectively responds to only one spatial frequency (the channel's *characteristic* spatial frequency). Suppose that in response to a visual stimulus, the output of each channel is a single number (or two numbers) proportional to the amplitude (or amplitude and phase) of the channel's characteristic spatial frequency in the stimulus. Suppose also that there exist a very large number of channels with characteristic frequencies covering the whole range of spatial frequencies responded to by the human visual system. Then, for any visual stimulus, the function relating the one output (or two outputs) of each channel to the channel's characteristic spatial frequency is a good approximation of the amplitude part (or amplitude and phase parts) of the Fourier transform of that stimulus. One might say, therefore, that such a set of channels actually performs a strict Fourier analysis. This extreme kind of spatial-frequency-channel model can be rejected easily, however, as a description of the bulk of psychophysical results.

Another extreme kind of "spatial-frequency-channel model" claims only that there are different subsystems sensitive to different ranges of spatial frequency. In this sense, spatial-frequency channels certainly do exist, for individual neurons are known to respond to different ranges of spatial frequency. (See Movshon, Thompson, & Tolhurst, 1978, for example.)

Investigators of visual psychophysics and perception, however, generally want something more of spatial-frequency channels. They want these channels to play an important role in explaining visual phenomena. Whether such interesting channels can reasonably be said to exist and, if so, what they are like are the questions of concern here.

Because the search for empirical evidence in favor of channels and the development of alternative conceptions of the channels have proceeded concurrently, further discussion of the concept of spatial-frequency channels is postponed until experimental results have been presented.

PSYCHOPHYSICAL EVIDENCE FOR SPATIAL-FREQUENCY CHANNELS

The use of spatial frequencies in describing visual stimuli is only two or three decades old, and the papers explicitly proposing spatial-frequency channels in the visual system are even more recent (Campbell & Robson, 1964, 1968; Enroth-Cugell & Robson, 1966). In the short time since then, however, a great deal of work has been done. I make no attempt, therefore, to mention all relevant work but give references to the original studies and to representative recent ones. This section of this chapter presents the major results from different kinds of psychophysical experiments and outlines the usual argument made from each kind for the existence of spatial-frequency channels. The next section describes tentative conclusions drawn from these experiments about the properties of the spatial-frequency channels.

Effect of Adaptation on Threshold Sensitivity and Suprathreshold Perceived Contrast

In these adaptation experiments, the observer inspects a suprathreshold grating (called the *adapting grating*) for several minutes and then looks at another grating (called the *test grating*). As Pantle and Sekuler (1968) and Blakemore and Campbell (1969) first showed, the contrast thresholds of test gratings close in spatial frequency to the adapting grating (and identical in all characteristics other than spatial frequency) are elevated after adaptation while the thresholds of test gratings further away in spatial frequency are unchanged. Similarly, the perceived contrast of suprathreshold test gratings close in spatial frequency to the adapting grating is reduced after adaptation, although the perceived contrast of test gratings further away in spatial frequency is unchanged (Blakemore, Muncey, & Ridley, 1971, 1973; Hertz, 1973).

These frequency-selective effects in adaptation experiments have often been explained as the result of "desensitization" (or "differential fatigue" or "adaptation") of multiple spatial-frequency channels. If an individual channel adapts (becomes less sensitive) after a period of being excited, the channels responsive to the adapting grating should be less sensitive after the observer inspects the adapting grating than they were before. The channels that are not responsive to the adapting grating should be unchanged by inspecting the adapting grating. On the assumption (1) that detection of a grating occurs whenever the response of at least one channel is big enough and (2) that perceived contrast of a grating is determined by the size of the response of the most responsive channel, this desensitization explanation qualitatively predicts the observed results. Accordingly, the properties of individual channels (in particular, absolute or relative bandwidths) have been inferred from selective adaptation results (e.g., Blakemore & Campbell, 1969; Graham, 1972). Although models of desensitization could be developed to predict quantitatively the effect, for example, of changing contrast in the adapting grating (analogous to Stiles' work on chromatic adaptation), very little theoretical work of this sort has been done (but see Dealy & Tolhurst, 1974, and Graham, 1970).

Recent explanations of selective adaptation effects have tended to invoke inhibition of one type or another rather than desensitization. Inhibition has been suggested by evidence that adapting to a combination of two frequencies produces less threshold elevation than adapting to one frequency alone (Nachmias, Sansbury, Vassilev, & Weber, 1973; Stecher, Sigel, & Lange, 1973; Tolhurst, 1972), by consideration of the minimum amount of contrast needed in the adapting grating (Dealy & Tolhurst, 1974), and by evidence that facilitation is observed when test and adapting frequencies are very far apart (DeValois, 1977; Tolhurst & Barfield, 1978).

If inhibition among channels is involved in the selective adaptation results, the properties of individual channels cannot easily be deduced from these results. It is not yet clear, however, what kind of inhibition, if any, is involved (Tolhurst & Barfield, 1978).

A few attempts have been made to extend these adaptation experiments to aperiodic stimuli like lines, edges, and random dots (Bagrash, 1973; DeValois & Switkes, 1978; Fiorentini, Sireteanu, & Spinelli, 1976; Georgeson & Sullivan, 1975; Legge, 1976; Sullivan, Georgeson, & Oatley, 1972; Williams & Wilson, 1978). If spatial-frequency channels are important, one would expect the effect of any adaptation stimulus on a subsequent test stimulus to depend on the degree to which the stimuli share spatial frequencies. To a large extent, the experimental results agree with this expectation. To compute exactly what effect one would expect, however, is complex. The aperiodic stimuli have all contained a wide range of spatial frequencies. The computation, therefore, would involve a large number of channels, and the exact mechanism of the selective adaptation effect is crucial. Little such computation has been attempted. In any case, although the experimental results are somewhat mixed, the evidence they give of the importance of spatial-frequency content seems quite good enough to encourage further exploration.

As far as one can tell at present, the results from selective adaptation experiments are consistent with the notion of spatial-frequency channels. Before taking these results as compelling evidence for the existence of such channels, however, one should ascertain that there is no other plausible explanation not involving such channels. One explanation often mentioned in this and similar contexts (Harris & Gibson, 1968) is local point-by-point adaptation. This kind of adaptation could occur as peripherally in the visual system as the retinal receptors and would not require any channels selective for spatial frequency. Results using stabilized-image techniques (Jones & Tulunay-Keesey, 1975) make this explanation unattractive, however. A second possibility is the model formulated by Wilson (1975). In this model, there is selective adaptation of the different lengths of connections between neurons. Since all the neurons have receptive fields of the same size, some people would not consider this to be a multiple-spatial-frequency-channels model. (See the section entitled "Summation at Threshold" for further discussion of receptive-field sizes and spatial-frequency channels.) This model is not a viable alternative either, however, because it cannot handle some of the experimental results described later. In short, because the available alternative explanations for spatial-frequency-selective adaptation effects can be ruled out, the adaptation effects do seem to be strong evidence for the existence of channels selectively sensitive to spatial frequency and active in psychophysical situations.

Effect of Masking on Threshold Sensitivity and Suprathreshold Perceived Contrast

In masking experiments, the detectability of one stimulus (the *test stimulus*) is measured in the presence of another (the *mask stimulus*). In some cases, the mask stimulus may precede or follow the test stimulus by a brief interval. (These cases are classified here as masking experiments rather than adaptation experiments because it seems likely, from what is known regarding the time courses of the responses of spatial-frequency channels, that the channels' responses to the mask stimulus and to the test stimulus do overlap in time.) When gratings are used in masking experiments, there is a frequency-selective effect at threshold resembling the effect in adaptation experiments; the threshold elevation is, in general, largest when the frequencies of test and mask gratings are identical. This is true whether the gratings overlap spatially (Legge, 1978; Mostafavi & Sakrison, 1976; Pantle, 1977; Sansbury, 1974; Stromeyer & Julesz, 1972; Tolhurst & Barfield, 1978) or do not overlap spatially (Rogowitz, 1977, 1978), and whether the mask grating is on continuously (e.g., Pantle, 1977) or occurs briefly at the same time as the test grating (e.g., Tolhurst & Barfield, 1978).

Masking also affects suprathreshold perceived contrast or clarity (Weisstein & Bisaha, 1972; Weisstein & Harris, 1980; Weisstein, Harris, Berbaum, Tangney, & Williams, 1977; White & Lorber, 1976). Very few frequencies of gratings have been used in these suprathreshold experiments, but the available evidence does show frequency-specific effects. Further, in Weisstein and coworkers' experiments with aperiodic stimuli, the amount of masking depends on the degree to which the spatial-frequency content of test and mask stimuli overlap.

These frequency-selective masking effects are usually interpreted as evidence for multiple spatial-frequency channels, presumably because it is difficult to think of any acceptable explanation not involving multiple channels. This prevalent interpretation seems reasonable, but how multiple channels lead to masking effects is still unclear. Masking effects, like adaptation effects, could be caused by (1) inhibition between the channels excited by the test stimulus and those exicted by the mask stimulus, or (2) desensitization of channels excited by the test stimulus because these channels are also excited by the mask stimulus. Masking effects could also be a result of some kind of interference when the mask stimulus and the test stimulus excite the same channels, or of difficulty at a higher-level decision stage in discriminating between the channels' responses to the test stimulus and those to the mask. Although the clearest statements of possible mechanisms for spatial-frequency-specific masking have usually postulated inhibition among channels (Breitmeyer & Ganz, 1976; Rogowitz, 1978; Weisstein, Ozog, & Szoc, 1975), the way in which this inhibition might work has not been thoroughly specified, nor have other explanations been ruled out.

Several results that seem to challenge any multiple-channels explanation of masking are worth mentioning briefly here. Masking occurs at periodicity frequencies not actually present in the stimuli (Henning, Hertz, & Broadbent, 1975); facilitation occurs for some combinations of mask and test frequencies (Nachmias & Sansbury, 1974; Stromeyer & Klein, 1974; Tolhurst & Barfield, 1978); and there are phase-specific changes in the amount of facilitation between a frequency and its third harmonic after adaptation to a square wave containing both (Sansbury, Distelhorst, & Moore, 1978).

Effect of Adaptation and Masking on Perceived Spatial Frequency

The experiments described earlier did not investigate directly the perception of spatial frequency or of size. They investigated directly the perception of contrast—either the perceived contrast of suprathreshold stimuli or the physical contrast of threshold stimuli. When the perceived frequency of test gratings is itself measured, it too is affected by previous exposure to an adapting grating (Blakemore, Nachmias, & Sutton, 1970; Blakemore & Sutton, 1969) or by the simultaneous presence of a spatially nonoverlapping masking grating (Klein, Stromeyer, & Ganz, 1974). Test gratings of frequencies somewhat higher than the adapting or masking grating look even higher; test gratings of somewhat lower frequencies look even lower; and test gratings of the same or very different frequencies seem unchanged.

The original explanation of this shift in perceived spatial frequency (Blakemore & Sutton, 1969) borrows two of the assumptions used in the desensitization explanations for the effects on threshold sensitivity and suprathreshold perceived contrast: (1) that multiple spatial-frequency channels exist, and (2) that adaptation to a pattern depresses the sensitivity of channels responding to that pattern. A third assumption, about the determinant of perceived frequency, is also needed to explain the shift in perceived frequency. (It takes the place of the assumptions about the determinants of contrast threshold and perceived contrast that were used in the earlier explanations.) Blakemore and Sutton (1969) assumed that perceived frequency is equal to the best frequency of whichever channel is most sensitive to the test grating (or perhaps to the frequency that equals some measure of central tendency of the distribution of all responding channels). According to these assumptions, adaptation depresses the sensitivity of the channels that responded to the adapting grating. Then, if the test frequency is close to the adapting frequency, the distribution of responses to that test frequency over all the channels will be different from usual, because the adapting grating will have desensitized some of the channels. The perceived frequency, therefore, will be biased away from the adapting frequency.

Unlike the corresponding assumptions about threshold and perceived contrast, this assumption about the determinant of perceived frequency requires that the output signals from different channels *not be* completely interchangeable. Any channel at all can be made to signal a threshold response or signal any given level of perceived contrast (by adjusting the contrast of the stimuli appropriately). But there is no way, according to this assumption about perceived frequency, that channels having different best frequencies can signal the same perceived frequency. Channels having best frequencies near a low spatial frequency, for example, cannot signal the same perceived frequency as channels having best frequencies near a high spatial frequency. In other words, to explain the shift in perceived spatial frequency, the channels are assumed not only to be selectively *sensitive* to spatial frequency but also to selectively *signal* spatial frequency.

Although the foregoing explanation of shifts in perceived spatial frequency is satisfactory at a qualitative level, there is some question as to its ability to account quantitatively for both threshold elevations and perceived spatial-frequency shifts simultaneously (Klein et al., 1974). Surprisingly little work has been published about the shifts in perceived spatial frequency, however, so one cannot be sure of the data, much less of the explanation.

Contingent Aftereffects of Adaptation

After an observer has adapted alternately to a red-and-black grating of one spatial frequency and a green-and-black grating of another spatial frequency, the appearance of a test grating that looked white and black previous to adaptation depends on its spatial frequency. (All gratings are of the same orientation.) If the frequency of the test grating is the same as one of the adapting gratings (e.g., the red-and-black adapting grating), the previously white bars will now appear faintly colored with a hue roughly complementary to that in the adapting grating (e.g., a faint green). If the frequency is slightly different from the adapting frequency, the colored aftereffect will be even fainter. These colored aftereffects contingent on spatial frequency, which are analogous to the original colored aftereffects contingent on orientation demonstrated by McCollough (1965), have been demonstrated in a number of studies (Green, Corwin, & Zemon, 1976; Harris, 1970, 1971; Lovegrove & Over, 1972; May & Matteson, 1976; Stromeyer, 1972; Teft & Clark, 1968). Contingent aftereffects can also be produced using orientation and spatial frequency, with the perceived orientation contingent on the stimulus spatial frequency or vice versa (Wyatt, 1974). Colored aftereffects contingent jointly on orientation and spatial frequency have also been reported (Wyatt, 1974).

These contingent aftereffects are sometimes qualitatively explained in a fashion similar to that used for explaining the perceived shifts in spatial frequency. The explanation invokes multiple channels selectively sensitive to narrow ranges of color, orientation, and spatial frequency (see Sigel & Nachmias, 1975, or Wyatt, 1974, for example). The channels are assumed to signal values selectively along each dimension. For example, the perceived value along a given dimension might be determined by the channel responding most to the stimulus or by some measure of central tendency of the distribution of responding channels. Adaptation is assumed to desensitize the channels responding to the adapting grating. Therefore, after adaptation to a red-and-black, low-spatial-frequency grating, for example, the channels sensitive to "red" wavelengths and low spatial frequencies will be fatigued. A white-and-black, low-spatial-frequency test grating will elicit less response from those channels sensitive to "red" wavelengths and low spatial frequencies than it usually does. Thus the perceived color will be biased away from red toward the complement of red, and white bars of the test grating will look faintly green.

These contingent aftereffects have certain properties that, to some people at least, seem odd if the aftereffects come from desensitization of sensory channels (Mayhew & Anstis, 1972). For example, contingent aftereffects last for weeks or even months (Jones & Holding, 1975; Riggs, White, & Eimas, 1974), and there is a report that they can be produced by imagining the color in the adapting stimuli (Finke & Schmidt, 1977). These properties have led to the suggestion that the contingent aftereffects are examples of learning, with the frequent implication that they may not demonstrate the properties of sensory channels at all (Harris, 1979; Mayhew & Anstis, 1972; Murch, 1976; Skowbo, Timney, Gentry & Morant, 1975). There is something appealing about the suggestion that "learning" or long-term changes in the nervous system are involved, but even if the effects are examples of learning, they may well demonstrate the properties of sensory channels. In order to learn, a learner must have input of some sort, and the form of that input must place restrictions on the form of learning.

Summation at Threshold

The experiments described so far all involved adaptation or masking. The experiments described next, however, do not. Rather, with the observers in what is assumed to be a neutral state of pattern adaptation, the detectability of compound patterns containing two or more component patterns is compared to the detectability of each of the component patterns alone. Underlying these experiments is the assumption that when two components excite the same channel, the results should show substantial "summation." In other words, a compound pattern containing two or more component patterns that excite the same channel should be a good deal more detectable than either component alone (because a single channel should be responding more to the compound than to either component). When the components excite different channels, however, the compound should be little or no more detectable than the most detectable component, because no channel responds more to the compound than to one of the components. (There may be a slight increase in detectability due to "probability summation"—that is, to the fact that two or more groups of channels may each have an independent chance to detect the compound pattern, whereas only one group of channels has a chance to detect a component by itself.)

Of all the kinds of experiments on spatial-frequency channels, these experiments exploring summation at threshold are the most extensive and the most quantitative. It is impossible to describe and review them adequately here. I have elsewhere (Graham, 1979) provided an introduction to these experiments and to the models describing them. For our present purposes, the most important summation-at-threshold experiments are those using combinations of sinusoidal components. These "sine-plus-sine" experiments have provided some of the strongest support for the notion of multiple spatial-frequency channels. A compound grating containing two or three sinusoidal components of very different frequencies is only slightly more detectable than its most detectable component; further, the detectability of such a compound grating does not depend on the relative phase of its components (Campbell & Robson, 1964, 1968; Robson & Graham in Graham, 1979; Graham & Nachmias, 1971; Graham, Robson, & Nachmias, 1978; Kulikowski & King-Smith, 1973; Lange, Sigel, & Stecher, 1973; Mostafavi & Sakrison, 1976; Pantle, 1973; Quick & Reichert, 1975; Sachs, Nachmias, & Robson, 1971). Such results are inconsistent with large classes of models postulating only a single channel but are consistent with models postulating multiple channels, each sensitive to a different range of spatial frequency.

Conceptual and experimental progress is being made on the specification of the exact properties of these channels. As several multiple-channel models have been refuted by summation-at-threshold experiments, the class of acceptable models has become better defined. The model that is emerging is usually described in terms of neural receptive fields, much as Thomas (1970) did (although a commitment to particular physiology is often not implied). In this description, each spatial-frequency channel contains a collection of "neural units" (or "neurons" or "detectors"), all with receptive fields of the same size. A collection of neural units with receptive fields having large excitatory centers is part of a low-spatial-frequency channel; a collection with receptive fields having smaller excitatory centers is part of a higher-spatial-frequency channel, and so forth. In one version of this model, the response of any one neural unit to a given pattern is variable, and the variability in different neural units is uncorrelated. An observer is assumed to detect a pattern whenever the response of at least one neural unit exceeds some threshold value. With these models, the detection thresholds of a large variety of patterns (periodic, aperiodic, and combinations thereof) can be quantitatively accounted for (for example, Graham, 1977; Quick, Mullins, & Reichert, 1978; Robson & Graham, 1978; Wilson & Bergen, 1977).

It is important to remember that at best, these summation-at-threshold experiments tell us about the relative sensitivity of different channels to different stimuli. The only signaling assumption required to explain these experiments is a simple one like: "An observer detects a pattern whenever at least one channel does." For these experiments, there is no need to assume that the output signals from different channels have qualitatively different perceptual consequences (much less to specify what the perceptual consequences actually are) and, therefore, no need to assume that the observer can even tell the signals from different channels apart.

Other Kinds of Psychophysical Experiments

The results of several other kinds of psychophysical experiments have been interpreted in terms of spatial-frequency channels. If all psychophysical experiments could be arranged on a continuum from "sensory" through "perceptual" to "cognitive," the ones described next would be further from the sensory end than the ones already described. An attempt to explain perceptual phenomena with the help of spatial-frequency analyses might well begin with a careful investigation of some of the results described next.

Uncertainty Effects. An observer's ability to detect a grating seems to depend on whether or not he or she expects a grating at that spatial frequency. More specifically, it depends on whether gratings of only one frequency are presented in a block of two-alternative, forced-choice trials or whether gratings of several frequencies are presented in random order on different trials of a block. The observer will be correct more often when the frequency is presented alone in a block than when it is intermixed with others (Graham et al., 1978).

Analogous uncertainty effects in the auditory domain have been interpreted by hypothesizing that the observer can pay attention to channels tuned to the appropriate frequency and thus avoid "false alarms" due to noise in other channels (Green & Swets, 1966). When frequencies are intermixed in a block, the observer must pay attention to a greater number of channels. One form of auditory experiment used unbalanced, intermixed blocks in which stimuli of one frequency, the "primary" frequency, were much more frequent than stimuli of any other frequency (Greenberg & Larkin, 1968; Macmillan & Schwartz, 1975). During such blocks, the observer would be expected to pay most attention to the channel tuned to the primary frequency. Then, detectability of stimuli that are processed by the same channels as the primary frequency should be less affected by the uncertainty of the intermixed condition than detectability of stimuli outside those channels' range. As expected, the effect of uncertainty is frequency selective, with the effect being least (the difference between alone and intermixed blocks least) at the primary frequency and becoming greater as the stimulus frequency moves away from the primary frequency. Recently, this frequency selectivity has also been obtained with visual spatial frequency (Davis & Graham, 1979).

Summation of Suprathreshold Contrast. In an extension of the summation-at-threshold experiments to a suprathreshold case, the perceived "overall contrast" of compound gratings containing two sinusoidal components widely separated in frequency was compared to that of the components alone (Arend & Lange, 1978; Hamerly, Quick, & Reichert, 1977; Quick, Hamerly, & Reichert, 1976). The perceived contrast of the compound was independent of the relative phase of the components and equal to the sum of the perceived contrasts of the components presented alone. Because the perceived contrast of any sinusoidal component was a very nonlinear function of physical contrast, accelerating quickly near threshold, the perceived contrast of a near-threshold compound was almost completely determined by the component with the greater perceived contrast. These results can be explained by assuming that perceived overall contrast is determined by summing the outputs of multiple spatial-frequency channels, where the output of each channel is a nonlinear function of contrast. The results of a study using gratings containing 10 sinusoidal components can also be explained this way (Abel & Quick, 1978).

Recognition Near Threshold. Several interesting and suggestive aspects of the recognition of near-threshold grating patterns have recently been discovered. For near-threshold sinusoidal gratings, an observer can recognize which of two widely separated frequencies is being presented just as well as he can detect that any pattern is there at all (Furchner, Thomas, & Campbell, 1977; Nachmias & Weber, 1975). Further, there is a range of near-threshold contrasts for which—although the observers can easily tell that a compound grating contains two frequencies (they can discriminate the compound from either frequency alone)—he cannot tell the relative phase of the two frequencies in the compound (Nachmias & Weber, 1975).

These recognition results are easily explained by multiple spatial-frequency channels if the observer can tell which channels are responding to a stimulus but cannot tell the relative phase of frequencies responded to by different channels. It is difficult, however, to imagine a single-channel model that would account for these results, for in most such models, frequency would be recognized on the basis of the response magnitudes at different spatial positions. For example, if peak responses occurred every one-third of a degree, the frequency would be recognized as 3 cycles per degree. But if observers can keep track of response magnitudes at different positions, they could probably recognize phase as well as frequency, because changing the phase changes the pattern of response magnitudes.

If two frequencies are quite close, an observer cannot always recognize which of the two frequencies is being presented (Hirsch, 1977; Thomas & Barker, 1977). This is the expected result, as the two frequencies sometimes excite the same channel.

Masking the Recognizability of Faces. On the basis of the masking results described earlier, one might expect that adding spatial frequencies near the frequencies in a photograph would interfere with the perception of the photograph more than adding other frequencies. Masking patterns of different spatial-frequency content were added to a photograph of Abraham Lincoln (Harmon & Julesz, 1973). Inspection of the resulting pictures shows that the recognizability of a portrait is greatly impaired by added frequencies close to the frequencies in the portrait but is much less impaired by frequencies farther away. How the perception of faces might be described in terms of outputs from spatial-frequency channels (or in any terms whatsoever) is far from clear. This demonstration is consistent, however, with the notion that channels selectively sensitive to spatial frequency are involved in the recognizability of faces.

Unstable Appearance of Compound Patterns. The appearance of a compound grating composed of widely separated spatial frequencies fluctuates (Atkinson & Campbell, 1974; Campbell & Howell, 1972). Sometimes the observer sees one component; sometimes the other; sometimes both. If the component frequencies are close together, however, the appearance is stable. These effects are interpreted as the result of alternation between the outputs of different spatial-frequency channels.

Aftereffects Seen on Blank Fields. Although many effects of adapting to gratings have been extensively investigated (as was described earlier), the spatial-frequency analogue of ordinary colored afterimages has not. Relatively little attention has been paid to the perceptions of an observer presented with an unpatterned field after long exposure to an adapting grating. The perceptions of this sort that have been reported and also some experienced during inspection of the adapting grating seem much more complicated than colored afterimages (Georgeson, 1976a, 1976b). Calling them "psychophysical hallucinations," Georgeson attributes them to antagonism between groups of orientation and spatial-frequency channels (but see MacKay & MacKay, 1976).

Texture Matches. In an experiment analogous to color matching, observers were required to match the appearance of suprathreshold texture patterns with mixtures of a few "primary" sinusoidal gratings of different spatial frequencies (Richards & Polit, 1974). The observers seemed to need only four different frequencies to make these matches. By the same logic that is used in color matching, these results were interpreted as demonstrating the existence of four (or more) spatial-frequency channels. These channels are not only selectively sensitive to spatial frequency; they also selectively signal spatial frequency.

Perceived Similarity of Textures. In a related experiment, a set of 30 texture patterns made up of combinations of seven spatial frequencies was used (Harvey & Gervais, 1977). Observers judged the perceived similarity of pairs or triplets of these patterns. Multidimensional-scaling representation of these data showed that a three- or four-dimensional space could account for the structure of the similarity judgments, as is consistent with an explanation based on three or four spatial-frequency channels.

TENTATIVE CONCLUSIONS ABOUT THE PROPERTIES OF SPATIAL-FREQUENCY CHANNELS

A century from now, the theory of visual perception may not even mention spatial-frequency channels. Their importance has certainly not been demonstrated conclusively or even compellingly. Further, even if the best explanation of each of the various psychophysical results already described does turn out to be spatial-frequency channels, the channels that explain one result (threshold elevation due to adaptation, for example) will not necessarily be the same as those that explain another (recognition near threshold, for example), and the interrelationship of the different channels may be quite complicated. Still worse, none of the channels may be of any significance in explaining everyday visual perception.

Keeping these cautions in mind, I would like to describe tentatively the picture of spatial-frequency channels that is beginning to emerge from psychophysical studies, with an emphasis on the aspects that seem relevant to spatial-frequency analyses of perceptual phenomena. In presenting such a picture, I do, of course, use personal judgment in weighing conflicting results. Let me point out that my bias favors summation-at-threshold experiments because those are the ones most extensively and quantitatively studied.

This description of spatial-frequency channels starts with what little is known about their signaling characteristics and the rules for combining their outputs. This is followed by a discussion of what is known about their interactions with one another. Channels' sensitivity to different stimuli (including their bandwidths and temporal properties) is then described at some length. The section ends with a brief discussion of how many channels there might be.

Signaling Characteristics and Combination Rules

If spatial-frequency channels are to explain perceptual phenomena, it is essential to answer the question of the perceptual consequences of the output from these channels. One aspect of this general question is that of signal form. Can the output of a channel (at any moment) be represented by a single number giving the magnitude of response? Or should the output be thought of as a function giving a number for each spatial position? (The number for a particular position could be the magnitude of the response of the neural unit having the receptive field that is located at that position and of the appropriate size for the channel in question.) Or does the output from a channel have some totally different set of characteristics?

Another aspect of the question of perceptual consequences is the issue of combination rules. How do the outputs from different channels combine or cooperate to determine the observer's response? What happens to the outputs of the set of spatial-frequency channels, usually conceived of as being relatively "early" in the information-processing hierarchy or relatively "peripheral" in the nervous system, as they enter into complicated perceptual and cognitive operations? (The question of combination rules is at least partially confounded with the question of direct influence from one channel on another. In most cases, it will be difficult to decide whether an observed interaction occurs at the higher level embodied in the combination rules or at the level of the channels.)

Unfortunately, very little is known about the perceptual consequences of the output signals from the spatial-frequency channels. Several points are worth making, however.

In discussion of color channels, the observer is frequently assumed to see a particular color whenever a particular channel responds. In some opponent-color theories, for example, the observer sees amounts of red (or green), blue (or yellow), and white (or black) that are proportional to the magnitudes of positive (or negative) outputs from the red-green, blue-yellow, and white-black channels respectively. In other words, any patch of color is thought to be perceptually analyzable into the component perceptions of red (or green), blue (or yellow), and black (or white), and these component perceptions correspond to the outputs of different channels (Hurvich & Jameson, 1974). In the case of spatial-frequency channels, however, it would not be reasonable to assume that a person sees a grating of a particular spatial frequency whenever a particular channel responds. Such an assumption is clearly inconsistent with everyday perceptions; the world does not, in general, look like superimposed gratings. More promising signaling assumptions are available.

The case of spatial frequency in vision may be similar to the case of frequency in audition. A compound grating consisting of two widely separated spatial frequencies can be seen as the juxtaposition of two sinusoidal gratings of different spatial frequencies, just as an auditory stimulus consisting of two widely separated frequencies can be heard as a chord of two tones of different pitches. Landscapes do not look like superimposed gratings, however; nor do spoken sentences sound like series of overlapping chords.

The example of color channels illustrates another point. In addition to the three opponent channels already referred to, there are indubitably at least three pigments serving color vision. These pigments are, of course, selectively *sensitive* to wavelength. Further, they *signal* wavelength selectively in the sense that the output signal from one pigment is not completely interchangeable with the output signal of another pigment at some point upstream (for normal observers). The output signal from any one pigment, however, does not correspond to any one of the perceptual components of color. In short, according to modern theory, there are at least two levels of color channels; both levels are selectively sensitive to wavelength and selectively signal wavelength, but only the higher level produces output signals that correspond in any simple way to perceptual components (Hurvich & Jameson, 1974). There is every reason to suspect that the processing of spatial-frequency information is at least this complicated, for we do not even have a clear idea of what the perceptual components of pattern or spatial vision might be.

Interesting answers to these questions of spatial-frequency channels' signaling characteristics and combination rules could be provided by psychophysical results that imply selective signaling of spatial frequency. The perceived spatial-frequency shift, the contingent aftereffects, and the recognition-near-threshold results are examples of such results and do suggest strongly that selective signaling exists. Unfortunately, most results of this kind have not yet been thoroughly investigated or are equivocal in the sense of being difficult to explain quantitatively within a multiple-channels framework.

Signaling characteristics are mentioned again in the section on temporal sensitivity characteristics.

Interactions Among Channels

The contrast thresholds of observers in a neutral state of pattern adaptation are consistent with channels that are independent of one another (see Graham et al., 1978, for example). This independence is of two kinds. First, there is probabilistic independence. The variability in the response of a channel to repeated presentations of the same stimulus is completely uncorrelated with the variability in every other channel's response; or, as is often said, there is probability summation among channels. Second, there is noninteraction among the average responses of different channels. That is, the average response of a channel to a stimulus (averaged over repeated presentations of the same stimulus) depends only on that channel's sensitivity to the spatial frequencies contained in the stimulus and is not influenced by how much or how little any other channel responds to that stimulus.

Several results are not easy to explain if the channels are always independent. In particular, several of the masking and adaptation results described earlier, where one of the patterns involved is suprathreshold or the observer is not in a neutral state of pattern adaptation, seem inconsistent with independent channels. Consequently, various forms of nonindependence have been proposed. Inhibitory interactions are the current favorite. The exact nature of this inhibition has not been specified, however. To put it another way, it has not yet been demonstrated that inhibition among channels could rigorously account for even one of the results it is invoked to explain, much less that one kind of inhibition could account for several of the results. Such a demonstration would handsomely repay the work necessary to attempt it.

Sensitivity Characteristics

Several types of psychophysical results imply channels that are selectively sensitive to spatial frequency without, however, implying channels that selectively signal spatial frequency. Many of these experiments (e.g., summation-at-threshold, effect of masking, and adaptation on thresholds) have been done and done carefully. They give us a good deal of information about the sensitivity of spatial-frequency channels to different stimuli.

Bandwidth. How large a range of spatial frequencies does an individual channel respond to? The very smallest estimates of bandwidth come from the original interpretations of sine-plus-sine summation-at-threshold experiments (Kulikowski & King-Smith, 1973; Quick & Reichert, 1975; Sachs, Nachmias, & Robson, 1971), but even according to these estimates, the range of frequencies responded to by an individual channel was not very narrow. Recent interpretations of sine-plus-sine experiments, which take into account independent variability in the responses of neural units at different spatial positions (Graham &

Rogowitz, 1976; Mostafavi & Sakrison, 1976; Quick et al., 1978) and also the interpretations of adaptation, masking, and near-threshold recognition experiments (see references given earlier), suggest a medium bandwidth (perhaps an octave at half amplitude).

Suppose (as was suggested in the earlier section, "Summation at Threshold") that a channel is conceived of as a collection of neural units having receptive fields that are identical in all properties except visual field location. Then one can interpret the bandwidth of a channel in terms of these receptive-field properties. Each receptive field has an excitatory center and inhibitory flanks and may have further secondary excitatory and inhibitory sections. The frequency responded to maximally by a channel is determined by the size of the receptive-field sections, with larger sections corresponding to lower frequencies. The bandwidth of a channel is determined by the number of sections in each receptive field. A channel in which each receptive field has only an excitatory center and inhibitory flanks responds to a relatively wide range of spatial frequencies. As secondary excitatory and inhibitory flanks are added to the outside of each receptive field, the range of frequencies to which the channel responds shrinks. In general, the greater the number of subsections in the receptive field, the smaller the range of frequencies; for if there are a very large number of subsections, even a small change in stimulus frequency away from the channel's best frequency produces a mismatch between the stimulus and the receptive field. (See Graham, 1979, for further explanation.) A bandwidth of about an octave corresponds to receptive fields that have an excitatory center, inhibitory flanks, and, perhaps, secondary excitatory outer flanks.

For a complete description of the sensitivity characteristics of a channel, one needs to know the orientation bandwidth as well as the spatial-frequency bandwidth; in other words, one needs to know the region in two-dimensional Fourier space to which the channel is sensitive. In terms of receptive fields, one needs to know how elongated the excitatory and inhibitory regions of the field are. If they are very long relative to their width, the receptive field will respond only to a very limited range of orientations. At the other extreme, if the excitatory and inhibitory regions are circular and concentric, the receptive field will respond equally well to all orientations. Much psychophysical evidence suggests that the channels respond only to a limited range of orientations. The narrowest orientation bandwidth estimates come from the original interpretations of summation-at-threshold experiments. More recent interpretations of these experiments and also of the results of selective adaptation and masking experiments suggest a rather broader bandwidth. (Many references are relevant here, but for lack of space, only a few representative ones can be given: Blakemore, Carpenter, & Georgeson, 1970; Blakemore, Muncey, & Ridley, 1973; Blakemore & Nachmias, 1971; Campbell & Kulikowski, 1966; Ellis, 1977; Kulikowski, Abadi, & King-Smith, 1973; Thomas & Shimamura, 1975; Tolhurst & Thompson, 1975.)

In any case, an individual channel certainly responds to more than a single spatial frequency and to more than a single orientation. The set of spatial-frequency channels is certainly not doing a strict Fourier analysis in the sense of the extreme model presented in the introduction.

Nonuniformity of the Visual Field. The visibility of high spatial frequencies relative to low ones decreases as one goes from the center to the periphery of the visual field. To account for this, the relative sensitivities of different spatial-frequency channels must shift as one goes from the center to the periphery. The quantitative description of this inhomogeneity is beginning to be worked out (e.g., Limb & Rubinstein, 1977; Robson & Graham, 1978; Wilson & Bergen, 1977; Wilson & Giese, 1977). This nonuniformity of the visual field is not so great, however, as to separate spatially the channels sensitive to different frequencies. At each location in the visual field, there must be more than one responsive spatial-frequency channel or, in other terms, more than one size of receptive field (Graham et al., 1978).

Temporal Sensitivity Characteristics

Fast and Slow Responses. Trade-offs between temporal and spatial visual resolution have long been noted. This spatiotemporal trade-off is evident in the different temporal characteristics of the visual system's responses to different spatial frequencies. The responses to low spatial frequencies are, in general, "faster" than the responses to high spatial frequencies. Let me distinguish among four possible ways in which responses might be "fast" or "slow" before discussing the temporal characteristics of different spatial-frequency channels.

Three of the ways concern responses to stationary stimuli. Consider the response to the onset of a stationary, long-duration stimulus. First, there may be a *delay* between the onset of the stimulus and the onset of the response. Second, there will be a *rise time* during which the response goes from zero to its peak. Third, after reaching its peak, the response may continue at the peak level, in which case we will call the response perfectly *sustained;* it may decay back to zero, in which case we will call it perfectly *transient;* or it may decay back to an intermediate level, in which case we will call it relatively sustained or relatively transient depending on whether the intermediate level is close to the peak level or to zero. In general, responses having shorter delay times, shorter rise times, or more transient time courses might be called "faster." Analogous distinctions can be made about the responses to the offsets of stationary stimuli.

If the system under consideration is linear or approximately linear, there is another way of describing the preceding three senses of "fast." Instead of describing the system's responses to stationary stimuli, one can describe its responses to stimuli flickering or drifting at different temporal frequencies. Delay time to a stationary stimulus translates into *phases* of responses to different temporal frequencies. Rise time translates into the *high-temporal-frequency cutoff*. The shorter the rise time, the higher the temporal frequencies to which the system can respond. Finally, a perfectly sustained response to a stationary stimulus corresponds to a complete absence of *low-temporal-frequency decline* (i.e., of decline in the system's sensitivity as one goes from medium to low temporal frequencies). In this case, the system responds equally well to all low frequencies. The more transient the response becomes, the more pronounced the low-frequency decline.

A fourth possible sense of "fast" depends on the responses to movement. The higher the rate of movement a channel is sensitive to, the faster the channel's responses might be said to be. If a channel is a linear system, its sensitivity to moving stimuli will be predictable from its temporal-frequency characteristics, and this fourth sense of "fast" will be closely related to the second—that is, to the rise time or high-temporal-frequency cutoff. The channel may not be linear, however, if, for example, it is truly directionally selective in the sense of responding best to one particular direction of motion, whatever the polarity of contrast in the moving stimulus. (See discussion in King-Smith & Kulikowski, 1975.) In that case, this fourth sense of fast is not equivalent to the first three.

Psychophysical Evidence for Spatiotemporal Interaction. A priori, these four meanings of "fast" do not necessarily go together. A channel may have, for example, a fast delay time and a slow rise time. As far as can be told from the limited psychophysical evidence available, however, channels sensitive to low spatial frequencies are faster in all four ways than channels sensitive to high spatial frequencies. Let's briefly review that evidence.

Both the longer reaction times to high-spatial-frequency gratings than to low- (Breitmeyer, 1975; Lupp, Hauske, & Wolf, 1976; and Vassilev & Mitov, 1976) and the precise timing of metacontrast effects using sinusoidal gratings as stimuli (Rogowitz, 1977, 1978) suggest that the time to the peak of a channel's response (delay time plus rise time) gets longer as the characteristic frequency of the channel gets higher.

Measurements of sensitivity to sinusoidal gratings presented with various time courses indicate that as the spatial frequency gets higher, the rise time of the response becomes longer (the high-temporal-frequency cutoff moves to lower temporal frequencies), and the response becomes less transient (the low-temporal-frequency decline becomes less pronounced). Sensitivity has been measured for gratings flickering or drifting at different rates (e.g., Robson, 1966; Van Nes, Koenderink, Nas, & Bouman, 1967; Watanabe, Mori, Nagata, & Hiwatashi, 1968), for stationary gratings exposed for different amounts of time (e.g., Legge, 1978; Nachmias, 1967; Schober & Hilz, 1965; Tynan & Sekuler, 1974), for gratings exposed twice with various interstimulus intervals (Breitmeyer & Ganz, 1977; Watson & Nachmias, 1977), and for a long exposure of a grating combined with a short exposure at various onset asynchronies (Tolhurst, 1975b). Also, the varying shapes of reaction-time distributions to gratings of different spatial frequencies suggest that the response to low-spatial-frequency gratings is more transient than the response to high (Tolhurst, 1975a). A similar conclusion may be drawn from studies of stimuli with specially designed time courses (Breitmeyer & Julesz, 1975; Wilson, 1978; and Wilson & Bergen, 1977).

For a discussion of evidence relevant to the fourth sense of "fast," which depends on the responses to moving stimuli, see Sekuler and Levinson (1977) and Sekuler, Pantle, and Levinson (1978). MacLeod (1978) reviews recent physiological as well as psychophysical evidence on spatiotemporal interaction.

Are There Only Two Kinds of Temporal Channels? As we have seen, there is an interaction between spatial and temporal visual sensitivity. Within a multiple-channels framework, this interaction can be expressed by saying that the low-spatial-frequency channels tend to be faster than the high ones. The details of this interaction are far from clear, however. Let me describe two rather different possibilities, neither of which I would be comfortable rejecting on the basis of current evidence.

One possibility is that there are only two kinds of temporal characteristics an individual channel can have. A channel's temporal characteristics can either be slow (sustained) or fast (transient). At the lowest spatial frequencies, there are only fast channels. At the highest spatial frequencies, there are only slow channels. For a wide range of intermediate spatial frequencies, there exist both slow and fast channels sensitive to each spatial frequency. As one changes spatial frequency in this range, the relative sensitivities of the slow and fast channels change dramatically (and their temporal characteristics may change a little). I am not sure that any investigators have proposed exactly this scheme, but suggestions of Kulikowski and Tolhurst (1973), Breitmeyer and Ganz (1976), Rogowitz (1977, 1978), and Watson (1977, 1978) are of this general sort.

The second possibility is that there is a continuum of kinds of temporal characteristics an individual channel can have, ranging from slowest (perfect sustained) to fastest (perfectly transient). At each spatial frequency, there is only one channel. The temporal characteristics change continuously from fastest to slowest as spatial frequency changes from lowest to highest.

Although these two possibilities seem quite different from each other, they both can account for the gradual change of temporal characteristics measured psychophysically as spatial frequency is changed. In fact, remarkably little evidence exists favoring one or the other possibility. One basic difference between the two possibilities is that the first but not the second postulates the existence of two distinct channels having the same preferred spatial frequency but different temporal characteristics. Watson (1977, 1978) found some evidence for the existence of two temporal channels at a single spatial frequency by measuring sensitivity to compound stimuli containing two flickering sinusoidal components of different temporal frequencies but the same spatial frequency. Further, some results from metacontrast experiments with pieces of sinusoidal gratings as stimuli are difficult to interpret unless there are two different temporal channels at a single spatial frequency (Rogowitz, 1977, 1978).

Another difference between the two possibilities is that the first, but not the second, divides the whole set of channels into two distinct groups on the basis of temporal characteristics and therefore suggests that these two groups serve two distinct functions. Perhaps (e.g., Breitmeyer & Ganz, 1976; Kulikowski & Tolhurst, 1973; Tolhurst, 1973) the set of faster, low-spatial-frequency channels is specialized for detecting temporal changes (as in motion perception, control of eye movements, and global processing of new scenes), whereas the set of slower, high-spatial-frequency channels is specialized for detecting spatial changes (as in pattern recognition, form perception, and local scrutiny of scenes).

If this distinction in function is valid, these two different groups of channels may have different signaling characteristics—that is, produce different perceptions. One difference between two perceptions is often taken as support for the notion of two distinct groups of channels. Observers are said to be able to distinguish between a perception of "movement" or "flicker" and a perception of "pattern," and to be able to set a threshold corresponding to either (e.g., Hood, 1973; Keesey, 1972; Van Nes et al., 1967; Watanabe et al., 1968). Assuming that these two different perceptions indicate which of the two groups of channels—fast or slow—is determining threshold, several investigators have used the two different thresholds to explore the spatiotemporal sensitivity of each group of channels (King-Smith & Kulikowski, 1975; Kulikowski & Tolhurst, 1973; Tolhurst, Sharpe, & Hart, 1973). This distinction between two different perceptions near threshold may correspond to other distinctions between perceptions of movement that have also led people to postulate two subsystems—one for handling movement irrespective of form, and one in which form is involved (Pantle & Picciano, 1976; Rashbass, 1968; Saucer, 1954).

The three pieces of evidence just summarized seem to favor the first possibility described—a dichotomy between slow and fast channels—over the second possibility—a continuum of temporal characteristics. Caution should be maintained, however, as the temporal summation effects obtained by Watson are small; the metacontrast timing effects obtained by Rogowitz are small; the phenomenological distinction between the two perceptions is not always easy for an observer to make and does not always lead to the expected results (e.g., Watson & Nachmias, 1977); and there are many other possibilities in addition to the two discussed.

Number of Spatial-Frequency Channels

The smallest number of spatial-frequency channels estimated is three or four on the basis of texture-matching and texture-similarity judgments (Harvey & Gervais, 1977; Richards & Polit, 1974). Wilson and Bergen (1977) also argue that four channels (four different sizes of receptive fields) at each location in the visual field are sufficient to account for an observer's sensitivity to a large variety of patterns. Four receptive-field sizes at each location, of course, imply many more altogether because of the nonuniformity of the visual field.

Although logically there is no necessary reciprocal relationship between bandwidth and number of channels, one might argue that the recent relatively broad estimates of bandwidth mean that relatively few channels would be sufficient to cover the range of visible spatial frequencies at any location in the visual field.

Early interpretations of selective adaptation effects suggested a rather large number of channels, because the largest threshold elevation was almost always at a test frequency identical to the adapting frequency (Blakemore & Campbell, 1969, for example). Also, early interpretations of the fact that the psychophysical contrast sensitivity function (which gives the observer's sensitivity to sinusoidal gratings as a function of spatial frequency) did not show individual bumps corresponding to individual channels suggested a rather large number of channels. When allowance is made for visual field nonuniformity and probability summation among channels, however, these two observations become consistent with a small number of receptive-field sizes at each spatial location.

On one hand, therefore, there may only be three or four spatial-frequency channels at each location. On the other hand, all the available evidence is probably also consistent with many more. Good estimates of the number of spatial-frequency channels remain remarkably elusive.

A NOTE ON SPATIAL-FREQUENCY CHANNELS AND PERCEPTION

The possible usefulness of spatial-frequency analyses in explaining visual perception has been pointed out by a number of investigators. Ginsberg (1971) concentrated on some of the patterns used to demonstrate Gestalt phenomena. He suggested that the stimulus information contained in certain ranges of spatial frequencies may be the basis for these phenomena. Many attempts have been made to explain visual illusions on the basis of spatial-frequency- and orientation-selective channels (e.g., Blakemore, Carpenter, & Georgeson, 1970; Bouma & Andriessen, 1970; Ginsburg, 1971; Oyama, 1977; Wallace, 1969). Georgeson and Sullivan (1975) suggested that spatial-frequency channels might be responsible for contrast constancy. The recognizability of visual stimuli despite changes in size and position has been attributed to encoding by spatial-frequency channels (e.g., Blakemore & Campbell, 1969). Numerous other suggestions have been made tentatively about the possible use of spatial-frequency channels in detecting Gibsonian texture gradients, in recognizing patterns (in letter confusion, for example), in encoding visual information in distributed form for memory storage (Weisstein & Harris, 1980), and so forth.

Of particular relevance to perceptual organization (e.g., Chapter 6 by Pomerantz in this volume) is the proposal that the low spatial frequencies in a pattern—which are processed by faster channels—are important for "global" processing, whereas the higher spatial frequencies—which are processed by slower channels—are important for "local" processing (e.g., Breitmeyer & Ganz, 1976; Broadbent, 1975, 1977; Kinchla, 1977).

Another application of spatial-frequency channels comes from work on "effortless texture discrimination," which is described in Julesz's Chapter 2 in the current volume. Julesz and his colleagues have studied visual textures to discover the conditions under which differences between regions of different textures become immediately and effortlessly apparent to observers. According to one form of Julesz's original conjecture, effortless texture discrimination is primarily based on the power spectra differences of textures and ignores their phase spectra. The power spectrum of a pattern is just the square of the amplitude part of a (two-dimensional) Fourier transform of the texture pattern; this amplitude part is, in turn, simply the function telling how much of each spatial frequency at each orientation is present in the pattern. Thus, two patterns differ in their power spectra if and only if differences exist between the amplitudes of corresponding sinusoidal components in the two patterns. Therefore, this form of the original conjecture can be reworded to say: If two texture regions differ sufficiently in the amplitudes of corresponding sinusoidal components in the two regions, the regions will be effortlessly seen as two separate regions. If two regions differ only in the phases of their components and not in the amplitudes, however, then the regions will blend into each other, appearing as one region of uniform texture.

Spatial-frequency channels could easily mediate the discrimination postulated by this conjecture if the channels signaled amplitude but not phase information to the higher centers. Then comparing the outputs of the channels in one region to the outputs in another region would reveal amplitude (power spectrum) differences between the two regions but not phase differences. (Because the spatial-frequency channels have greater-than-zero bandwidths, their outputs are not precisely proportional to the amplitudes of individual sinusoidal components. That is, the channels do not perform a strict Fourier analysis. They may perform one good enough to account for texture discrimination, however, for no one claims that the quoted conjecture is precisely true.)

These explanations of perceptual phenomena in terms of spatial-frequency analyses are in various states of development. Some have been investigated extensively; some barely at all. No one of the suggestions, however, has been developed enough to be compelling. This state of affairs should come as no surprise. The problems of explaining perceptual phenomena are much too difficult to permit solutions without an enormous amount of work.

The available psychophysical evidence clearly demonstrates that there are channels (or subsystems or mechanisms) in the visual system that are selectively *sensitive* to spatial frequency. Although quantitative, rigorous explanations of many of the results remain to be constructed, a very large amount of data can be organized with the help of the concept of such channels. The psychophysical evidence also suggests that there are channels that selectively *signal* spatial frequency in the sense that the output of one channel is not completely interchangeable with the output of another. What happens to these outputs from the channels, however, is still unclear. We do not know how they combine with or inhibit one another. More generally, we know very little about their perceptual consequences. It is, of course, the nature of these perceptual consequences that will largely determine how useful the channels are in explaining perceptual phenomena.

ACKNOWLEDGMENTS

The preparation of this chapter was supported in part by a grant from the National Science Foundation (BNS 76-18839). I am grateful to E. Davis, B. A. Dosher, J. E. Hochberg, D. C. Hood, P. Kramer, C. R. McCauley, J. Nachmias, J. P. Thomas, and A. B. Watson for their helpful criticisms of early drafts of this chapter.