GRAHAM (IN PRESS.) International Encyclopedia of Neurology
Psychiatry, Psychoanalysis and Psychology. ed. B. Wolman

SPATIAL FREQUENCIES

The concept of "spatial frequencies" has become increasingly important in the study of vision during the past twenty years. This concept is based on the fact that any complex waveform can be analyzed into a number of elementary sinusoidal waveforms which, when added together, reproduce the original waveform. The mathematical techniques using this fact are known as Fourier analysis. They have been of use in audition, for example, since they permit the decomposition of complex sounds into their constituent simple harmonic variations in sound pressure level (pure tones). These same techniques have become very prominent in the study of visual processes in recent years (Cornsweet, 1970; Harris, 1975).

When these techniques are applied to vision, it is possible to analyze a visual scene into simple sinusoidal variations in luminance, which are related to the sizes of the elements comprising the scene. The scene itself is then the sum of these sinusoidal variations .

To fully understand these techniques of analysis of visual scenes it is easiest to begin by considering one very simple pattern, a *sinusoidal grating*. A sinusoidal grating is a pattern which looks like a blurry, alternating set of dark and light stripes. In such a pattern the luminance (light intensity) in the direction perpendicular to the stripes varies sinusoidally while the luminance in the direction parallel to the stripes remains constant.

Several parameters describing a sinusoidal grating are important. The one of most concern here is *spatial frequency*, that is, the number of cycles of sinusoid per unit distance in the pattern or, equivalently, the number of dark-bar-light-bar pairs per unit distance in the pattern. Spatial frequency and size of bar are inversely proportional to each other, as are the wavelength and frequencies of sound or radio waves. A grating with very narrow bars has a high spatial frequency, a grating with very wide bars has a low spatial frequency. The mean luminance of the grating is the average of the luminances across the whole spatial extent of the grating. The contrast of the grating is some measure of the difference between the most luminous and least luminous points, usually taken to be one-half that difference divided by the mean luminance. One of the attractions of sinusoidal gratings to visual scientists has been that the mean luminance of a grating can easily be held constant (thus, presumably, keeping the observer in a 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. It can be proven, following the work of the eighteenth century mathematician Fourier, that any such stimulus could be constructed by adding together sinusoidal gratings of different frequencies, contrasts and phases (positions). Consider, for example, a set of alternating black and white stripes, or in other words, a periodic square wave grating in which one region of dim luminance is abruptly followed by a region of high luminance which, in turn, is abruptly followed by a region of dim luminance, etc. This square-wave grating can be described as the sum of a set of simple sinusoidal variations in luminance. This set includes one sinusoidal grating with a period or fundamental frequency identical to that of the original square wave pattern. Another component has a frequency three times that of the fundamental frequency of the pattern and of one third its amplitude. A complete analysis of the pattern will show that it contains an infinite number of sinusoids where all of them are odd harmonics of the fundamental frequency, and their amplitudes are inversely proporitonal to their order. (The third harmonic has one third the amplitude of the fundamental, the fifth harmonic has one fifth the amplitude of the fundamental, and so forth. It will be seen that the very high-frequency components of the square wave pattern become vanishingly small as their frequencies increase). Analyses of other patterns show that their compositions differ from that of the square wave pattern discussed here. Thus, any one-dimensional visual stimulus can be described as containing certain spatial frequencies - the spatial frequencies of the sinusoids which, when added together, form the original stimulus.

This discussion has been concerned with patterns which vary in one dimension. The analysis into spatial frequencies is not restricted to visual stimuli that vary along one dimension, although it is more difficult to describe analyses of two or three dimensional patterns (to say nothing of four dimensional analyses which involve time varying patterns).

Pictures of fine-grained sandpaper, for example, contain high spatial frequencies in two dimensions. Or, to put it another way, a fine-grained sandpaper pattern could be formed by adding up sinusoidal gratings of high spatial frequencies and of many orientations with respect to each other (stripes in the different grating would not be parallel to each other). Similarly, a course-grained sandpaper pattern contains lower spatial frequencies in both dimensions.

This analysis of visual patterns into spatial frequency components has become important in the study of vision because it is a natural and fruitful method to use if one is dealing with certain kinds of systems called "linear systems". Some early theories suggested that, at least to a first approximation, the human visual system might be a linear system. By definition, the response of a linear system to a stimulus which equals the sum of component stimuli is exactly equal to the sum of the responses to the component stimuli by themselves (that is, "the response to a sum equals the sum of the responses"). This is known as the superposition principle. Since any visual stimulus is equal to the sum of sinusoidal stimuli, the response of a linear system to any visual stimulus would equal the sum of the responses to the appropriate sinusoidal stimuli. If the visual system were linear and its responses to sinusoidal stimuli were known, its responses to any stimulus could be computed by analyzing the stimulus into its sinusoidal components. So using sinusoidal gratings and discussing all stimuli in terms of the spatial frequencies they contain is a natural way to study a visual system assumed to be linear.

The idea that the whole visual system might be described as a single linear system has become outmoded, for good reason. But talking about spatial frequencies has not. For one thing, sinusoidal gratings continue to be used in neurophysiological studies of isolated parts of the visual system. A part may well be linear even if the whole visual system is not. And secondly, spatial frequencies are much discussed in the context of "multiple-channels models" of the visual system.

In the "multiple-channels models", the visual system is described as containing multiple systems (or "channels") all of which simultaneously respond to a visual stimulus. In the version of multiple-channels model proposed by Campbell and Robson (1968), spatial frequencies played a particularly prominent role. These authors suggested that each of the channels responds only to patterns that contain spatial frequencies in certain narrow range. Also, each different channel responds to frequencies in a different narrow range. Some investigators have referred to these proposed channels as "spatial frequency channels."

Evidence from psychophysical experiments has been accumulating which is not inconsistent with such a multiple channels model. This evidence will be discussed below. And the eventual success or failure of the model, which describes the early stages of the whole human visual system, should probably depend on psychophysical evidence. But since the development and the study of the model have been much influenced by neurophysiological evidence, a brief discussion of the relevant neurophysiology is appropriate.

For several decades, it has been known that neurons in visual systems of all but the lowest animals are specialized. Different neurons require different stimuli for maximal response. Particularly important for the concept of spatial frequency channels is the fact that some neurons in the visual cortex of higher animals respond to bars of a particular size but not to bars much wider or narrower. (The bars must also be of a particular orientation, and, in fact, there is a corresponding notion of "orientation channels.") It is possible that these size-selective neurons form the psysiological substrate for spatial frequency channels. The neurons responsive to wide bars might form the low spatial frequency channels, those responsive to narrow bars might form the high spatial frequency channels, etc. It is not clear, however, due to insufficient neurophysiological and psychophysical evidence, whether or not the cortical neurons and hypothesized spatial frequency channels have corresponding properties. Perhaps in the next few years the appropriate experiments will be done.

Most of the psychophysical experiments that have produced evidence consistent with a model of multiple spatial frequency channels fall into one of three classes: (a) detection-summation experiments, (b) adaptation (or masking) experiments measuring threshold effects, and (c) adaptation (or masking) experiments measuring supra-threshold effects.

In detection-summation experiments, the detectability (visibility) of a compound pattern containing two or more sinusoidal components is compared to the detectability of each of the sinusoidal components alone. If the sinusoidal components are of very different spatial frequency, they should, according to the multiple-channels model, excite different channels. Then the response of each individual channel to the compound pattern should be no greater than its response to one of the component sinusoids. Suppose an observer can tell that a pattern is present whenever a single channel responds. Then, to a first approximation, the compound pattern will be no more detectable than the most detectable of its sinusoidal components. And, further, the relative phases of the component sinusoids should not matter. In the rather large number of experiments of this kind that have been done, the result predicted by the multiple-channels model has always been confirmed (e.g. Graham and Nachmias, 1971).

In adaptation experiments, mentioned in (b) and (c) above, the observer inspects a supra-threshold grating (called the adapting grating) for a relatively long period of time (some minutes) and then looks at another grating (called the test grating). 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. But the channels that are not responsive to the adapting grating should be unchanged by inspecting the adapting grating. Thus, the thresholds for test gratings of frequencies very close to the adapting grating's frequency should be elevated after the inspection period while the thresholds for test gratings of other frequencies should be unchanged. This frequency-selective threshold elevation has been overwhelmingly confirmed in a number of experiments.

A similar frequency-selective effect of adaptation on the appearance of supra-threshold test gratings is also predicted by a multiple channels model. One would expect the perceived appearance of a test grating having a spatial frequency very far from the adapting grating's frequency to be unchanged (because none of the channels responsive to the test grating will have been affected by the adapting grating). But the perceived appearance of a test grating close in frequency to the adapting grating might well change. The exact changes predicted in perceived appearance depend on the assumptions made about the relation between the output of the channels and perceived appearance. One possible change would be for a test grating slightly lower in frequency than the adapting grating to look lower still, and test gratings slightly higher in frequency to look higher still. In the original experimental reports, both the unchanged appearance of test gratings different in frequency from the adapting grating and also the changed appearance of test gratings close in frequency were reported. There has been some difficulty however, in replicating this result. It now seems as if there is indeed some change in perceived appearance but the change occurs even for test gratings quite far in frequency from the adapting grating as well as for gratings close in frequency. As will be discussed below, this difficulty is symptomatic of the current state of research in this area.

Masking experiments (mentioned parenthetically in (b) and (c) above) involve simultaneous rather than successive presentation of two gratings--a masking grating (analagous to the adapting grating) and a test grating. The experimental results and the logic relating these results to the multiple channels model are very closely analogous to those for the adaptation experiments described above.

All the experiments mentioned above made use of periodic patterns, either simple sinusoidal gratings or more complex gratings containing superimposed sinusoids. If the multiple channels model is correct, however, it ought also be able to explain the perception of aperiodic patterns like edges or lines.

(Aperiodic patterns are also subject to Fourier analysis and can in fact be reproduced by superimposing very large numbers of sinusoids.) Unfortunately, making predictions about aperiodic patterns from a multiple channels model is not a trivial task. Since any aperiodic pattern contains a broad range of spatial frequencies , many channels will be involved in the perception of such a pattern. Therefore, one must deal with questions concerning the number of channels and their bandwidths (the range of frequencies to which each responds) in order to make predictions about the perception of aperiodic patterns. Recent attempts to deal with these questions have been moderately successful (see Harris, 1975).

In conclusion, a model of multiple spatial frequency channels does predict results which approximate those actually obtained in a wide range of psychophysical experiments. However, the model cannot account for the fact that different classes of experiment appear to yield descriptions of the channels which are widely divergent. Thus, experiments of type b) suggest channel bandwidths that are much broader than those currently estimated from experiments of type a). Bandwidths estimated from experiments of type c) are broader still. Moreover, there has been little progress in extending the multiple channels model to supra-threshold perception. It is not clear whether these difficulties result from a major flaw in the multiple channels model or from insufficient ingenuity on the part of theoreticians.

References

Campbell, F.W., & Robson, J.G., Application of Fourier Analysis to the visibility of gratings, J. Physiol., 1968, 197, 551-66.

Cornsweet, T.N. Visual Perception, New York, Academic Press, 1970.

Graham, N., & Nachmias, J., Detection of grating patterns containing two spatial frequencies: A comparison of single-channel and multiple-channels models. Vision Research, 1971, 11, 251-59.

Harris, C.S. (ed.), Visual Coding and Adaptability, New Jersey, L. Erlbaum Associates, 1975.

Also see:

- Psychophysics
- Auditory Psychophysics
- Feature detection
- Sensation and Perception