What Psychophysical Thresholds Tell Us

The mapping from the physical world to consciousness can be described as a manifold which is smooth and continuous almost everywhere.

Robert E. Armstrong

The author is a student ostensibly with the psychology department of Northwestern University

When we open our eyes we see entire scenes, perceptions, not colored points, sense data, or sine functions. These perceptions in their entirety form the evidence from which psychophysicists must meld their theories. Essentially, psychophysics is the specification of a relation $F:S \rightarrow P$ which takes the space of all possible stimuli, S , into the space of all possible conscious perceptions, P . There are considerable limitations to what can be inferred about the details of neurophysiological organization on the basis of knowledge obtained about F . Even if we make the assumption that there is a one to one relationship between brain states and conscious states, we still are not much further within.

In particular, psychophysicists cannot resolve uniquely one of the age old thematic questions (1) in psychology - what are the elemental units of perception. During the last decade or so, a number of psychophysicists have proposed and claimed to find evidence for a variety of elements or features. These have included lines and edges (2) and sinusoids (3, 4, 5, 6).

If any of these functions do form the bases of our perceptual world, we probably cannot determine it by psychophysical methods. Scalpels and microelectrodes are probably the only tools which will yield unambiguous results (7).

In this paper, I will discuss in detail only the intrinsic ambiguity of the results of subthreshold summation experiments, however the implications for the general problem of interpretation facing psychophysicists will be mentioned. I will show that a great deal of psychophysical threshold data can be explained by rather naturalistic assumptions about what a subject tries to do in a threshold experiment. I will show that internal noise determines the extent to which the structure of the system is elucidatable by threshold experiments, and that Signal Detection Theory (8, 9) must be considered in the interpretation of almost all threshold data. More than anything else, I want to express the belief that the explanations offered for perceptual systems need to be couched in terms of concepts which allow manipulation of entire patterns of stimuli and of neural activity as single entities. These concepts are extant in modern control theory, systems theory, and dynamical systems theory (10). The basic concept is that an entire pattern can be envisioned as a point in a very high, or even infinite dimensional space. As the stimulus changes from pattern to pattern, the point specifying its perceptual representation moves on a very high dimensional surface or manifold determined by F . Psychophysical threshold tasks test the subject's ability to discriminate between points or between paths on this manifold. Distinguishability of points in the world is determined by the stochastic nature of the stimulus pattern itself and the stochastic nature of F . F is continuous and smooth (differentiable), almost everywhere, the exceptions being boundaries along which major changes or state take place. The differentiability combined with the statistical nature of F induces a Euclidean or more properly Reimannian metric on the space of patterns almost everywhere. This causes threshold loci to generally be elipsoidal.

These notions are elaborated below and several experimental consequences are pointed out.

Terminology

Hopefully this paper will be accessible to people with moderate levels of mathematical training. The major concepts required are those covered in an elementary linear algebra course and most of the mathematics is virtually identical to that of analysis of variance and factor analysis. The basic objects we will discuss are patterns. As I said earlier, we will conceive of these as points or vectors in multidimensional spaces. I will denote patterns (vectors) by lower case boldface letters, e.g. a, t, etc. In visual experiments, typical stimulus patterns are functions which assign a luminance level to each point in the visual field. A function is a vector in an infinite dimensional space. I will label entire functions with bold face letters just like any other vectors. The value of a function at a particular point will be denoted by a form like f(x,y). A particular element of a finite dimensional vector will be denoted by a subscripted letter, e.g. a_i. As long as the resulting pattern does not assign negative luminance values to any points, we can add or subtract any two patterns to get a third. We can also multiply a pattern by a constant. This simply changes the amplitude of the pattern. Constants will be denoted by ordinary letters.

We will be particularly concerned with functions, also called transformations or mappings, which change one pattern into another. I will denote transformations by capital letters, e.g. F, T, L, etc. A transformation, T is continuous near a point (pattern) a if the transformed pattern T (a) is close to T (a+ε) when ε is small so that a+ε is close to a.

To make sense out of the previous sentence, we need some notion of distance. We will assume Hilbert space or Euclidean structure on all our vector spaces. The analytic and empirical justification for this will be demonstrated later. At this point I will only comment that this structure is prerequisite for the notions of angle and orthogonality to make any sense. It is also requisite for the application of the Fourier transform. Hilbert space structure means that an inner product is defined on vectors in the space and distance is defined in terms of this inner product. Specifically, we will define the inner product of two finite, dimensional vectors to be

$$\langle \underline{x},\underline{y}\rangle = \sum_{i=1}^{n} x_i y_k$$

for infinite dimensional vectors we define

$$\langle \underline{x},\underline{y}\rangle = \int_R x(r)y(r)dr$$

where R is the region on which x and y are defined.

The Euclidean length or "norm" of a vector is defined as the square root of the inner product of the vector with itself. We denote the norm of the vector x as

$$||\underline{x}|| = \sqrt{\langle \underline{x},\underline{x}\rangle}$$

Two vectors are said to be orthogonal or perpendicular to each other if their inner product equals zero.

A unit vector is a vector whose length equals 1. Any vector can be made into a unit vector by dividing it by its length.

When manipulating linear transformations, standard linear algebraic notation will be used for both finite and infinite dimensional vectors and transformations.

The concept of feature detectors

Perhaps the strongest claim of this paper is that there are major limitations on what can be said about the detailed internal structure of the brain on the basis of psychophysical evidence. I will begin by examining the notion of "feature detectors" which has arisen in recent years.

The original impetus for the psychophysicists' search for feature detectors probably derives from Lettvin, Maturana, McCulloch, and Pitts classic 1959 paper, "What the frog's eye tells the frog's brain" (11), and seminal work of Hubel and Wiesel on the response of neurons in the visual cortex of the cat and monkey (12). These papers in turn were motivated by concepts arising from early efforts at mechanical pattern recognition.

Around the same time, methods of linear systems analysis were borrowed from certain fields of engineering and applied to the analysis and description of the performance of the visual system. Studies using these techniques are recognizable by the frequency of their use of repetetive one-dimensional patterns described as sums of sine and cosine functions. In 1968, Campbell and Robson in a highly influential paper (3) merged the notions of feature detection and Fourier analysis and proposed that the visual system was composed, at least in part, of channels or filters responsive to relatively narrow bands of spatial frequencies. A theory was developed which might be called the "multi-channel linear filter hypothesis," which proposed that the perception of patterns is mediated by the responses of channels or detectors which are maximally responsive to particular pattern elements and respond linearly, at least near threshold. A pattern was assumed to come to threshold when the response of some channel reached threshold. Campbell and Robson proposed that such a model could explain discrepancies between data they obtained and a simple "peak detector hypothesis" which proposed that the threshold for a pattern is determined simply by the peak to peak amplitude of the MTF weighted wave form. I will show later that neither of these two assumptions is very plausible from the standpoint of signal detection theory and that neither fits Campbell and Robson's own data.

However, let us stop for a moment and examine more closely what is meant by a feature detector. Apparently the outputs of these detectors are supposed to be the grounding for our perceptual experiences. Consider a simplified outline of what is known physiologically to happen in the early stages of the visual system. An external pattern $\underline{s}$, is transformed by the optics of the eye into a pattern of light, $T_o(\underline{s})$, on the retina which causes a pattern of activity $T_r(T_o(\underline{s}))$ in the receptors of the retina. This pattern is modified by the middle layers of the retina and produces a pattern of retinal ganglion cell output $T_g(T_r(T_o(\underline{s})))$. The next way station in the chain to cortex is the lateral geniculate. It produces a transformation in the incoming pattern $T_{lg}(T_g(T_r(T_o(\underline{s}))))$. Finally, this transformed pattern reaches visual cortex and we will say that at the first stage a transformation T_C is performed. We will lump everything else which happens to the pattern, from this point until it reaches consciousness, under T_Ω. I could further refine and dissect these transformations but it is not necessary for making my point. The overall function F, is of course simply equal to the composition of all of these functions

$$F = T_\Omega \circ T_C \circ T_{lg} \circ T_g \circ T_r \circ T_O$$

In some sense, the output of each of these stages is a grounding for our perceptual world. The first neural stage, $T_r \circ T_O$, provides a basis of colored points, analagous to a basis of Dirac δ-functions. $T_C \circ T_{lg} \circ T_g \circ T_r \circ T_O$ provides a basis of Hubel-Wiesel cells -- edges, bars, Lie derivatives or whatever. The question psychophysicists have tried to answer is, given F can the individual effects of T_r through T_C be made visible in any way? Psychophysicists are looking for some signature in our visual experience left by each of the stages. It could be argued that the best system would not leave any. The stages in a linear system do not.

In general, when we gain knowledge about the behavior of a black box or lumped system like F, we can specify its contents, its actual organization only down to some equivalance class of functions. For instance, the existence of series representations, Taylor, Fourier, Volterra, or whatever, for almost any arbitrary system demonstrates that the performance of a black box can be duplicated by organizations which are highly unlikely to bear much resemblance to the actual system being mimicked. Engineers usually are not concerned by their lack of uniqueness, but that is specifically what we are concerned with here.

In brief, all that psychophysics can specify is the equivalence class of structures which embody the transformation F. However, I will show that some interesting things can already be specified about the equivalence classes of F near specific points. For this we need some more theory.

A theory of thresholds

Thresholds for discriminating patterns are obviously determined by many factors including such things as abilities of attention, memory, fixation, and so forth. However, a fairly naturalistic model of what a subject does in an experiment combined with the notions about perceptual systems mentioned earlier can lead to some rather strong predictions about threshold behavior. Since a subject has similar information available no matter what the experimental paradigm, we will discuss just one which will be sufficient to get the major ideas involved across. The notions are essentially those of modern signal detection theory (SDT) as elaborated in electrical engineering and communication theory (8). Psychologists have for some time been acquainted with a subset of this theory (9) which is inadequate for description of vision experiments. Furthermore, some nonlinear transformations have been proposed in the name of psychological signal detection theory (PSDT) which will be explicitly rejected here (13).

We must start by considering the description of the stimuli. The stimulus patterns we will be concerned with are stochastic functions assigning each point in some region of the plane a luminance value, or more generally a distribution of intensities of the various wavelengths of light (14). So a pattern is a function $f(x,y)$. Since any physical pattern contains noise -- quantal if no other -- $\underline{f}$ is a stochastic function. If the noise is purely quantal, $f(x,y)$ has a Poisson distribution. For our purposes we will make the simplifying assumption that we can express $\underline{f}$ as the sum of its expected or average appearance $\underline{\bar{f}}$ and a noise component $\underline{e}$ where $e(x,y)$ is approximately normal with zero mean and covariance function $C(x,x',y,y')$. If the noise is Poisson, the noise is virtually independent from point to point and has variance equal to the mean luminance at that point, $f(x,y)$.

Notice that our stimuli are not expressed as functions of time. Since in this paper I am not going to discuss the dynamics of F but consider only the detection of perturbations from steady states, the time series character of the actual stimuli will be supressed. The justification for this will be presented elsewhere.

It will be useful to think of a pattern $\underline{f}$ as being an ensemble of patterns $\underline{f}_{\alpha}$ $\alpha=1,2,3,\ldots$ which on the average look like $\underline{\bar{f}}$ but each of which has had a pattern of noise, $\underline{e}_{\alpha}$, added to it. On any particular presentation of the putative stimulus $\underline{f}$, some element $\underline{f}_{\alpha}$, of the collection actually occurs. It will be assumed that through repeated exposure, the subject gains knowledge of the approximate appearance of $\underline{\bar{f}}$ and the approximate "shape" of the cloud of $\underline{f}_{\alpha}$'s which occur. Thus he develops an intuition as to the likelihood of deviations of various sizes in various directions in stimulus space.

Now assume that a subject is presented an adapting stimulus $\underline{a} = \underline{\bar{a}} + \underline{e}$, and a test stimulus which is a perturbation of $\underline{a}$. That is, $\underline{t} = \underline{\bar{a}} + \delta \underline{\bar{p}} + \underline{e}$, where δ is a small number and $\underline{\bar{p}}$ is a unit vector. We are assuming that the distribution of errors $\underline{e}$ around $\underline{\bar{t}}$, i.e., $\underline{\bar{a}} + \delta \underline{\bar{p}}$ is approximately the same as the distribution around $\underline{\bar{a}}$ alone. This is reasonable if the maximum deviation over all points (x,y) of $t(x,y)$ from $a(x,y)$ is only a few percent of the value of $\bar{a}(x,y)$.

The subject, of course, has no immediate knowledge of $\underline{t}$ or $\underline{a}$. The observer's definition of what $\underline{t}$ and $\underline{a}$ "really" are depends on correlations between the readings of various instruments which have in some sense better acuity than he does. The subject's perception of $\underline{t}$ is given by $F(\underline{t})$. Remember, our fundamental assumption stated earlier is that F is smooth, i.e., differentiable in the vicinity of $\underline{a}$. Specifically, this means that there exists a linear transformation L such that

$$F(\underline{\bar{t}}) = F(\underline{\bar{a}} + \delta \underline{\bar{p}}) = F(\underline{\bar{a}}) + L(\delta \underline{\bar{p}}) + \psi(\delta \underline{\bar{p}}) \qquad (1)$$

where $\psi(\delta \underline{p})$ goes to 0 as the "size" (norm) of $\delta \underline{p}$ goes to zero. (15)

In other words, a small change in the physical pattern produces a small change in the pattern perceived, twice as large a change in the physical stimulus produces about twice as large a change in the output, and the change produced by the sum of two perturbations is approximately equal to the sum of the changes produced by the two of them acting individually. That is, L is a function with the property that

$$L(a\underline{x}+b\underline{y}) = aL(\underline{x}) + L(\underline{y}) \quad (2)$$

Unfortunately, differentiability is an abstract concept with almost no empirical import. Note that (2) implies that $L(\delta \underline{p}) = \delta L\underline{p}$. For the real world, we require a somewhat stronger assumption – that F can be approximated by a linear function over the range of threshold size perturbations. This is not unreasonable since, for instance, the log function which is a reasonable approximation to overall visual system response is negligibly nonlinear even for perturbations as large as 10%. Most thresholds are achieved with smaller deviations than this and we will restrict our attention to these cases. It should be noted the transformation L (in general) depends upon the point $\underline{\bar{a}}$, i.e., the appropriate approximation to F is different for different values of $\underline{\bar{a}}$.

We will assume that the noise $\underline{e}$ is also small relative to $\underline{\bar{a}}$ which it must surely be since no threshold strategy could produce better data than that dictated by the intrinsic noise in the stimulus and we are concerned with situations in which threshold is reached with small perturbations of the stimulus. Therefore, we can write

$$\underline{t}^{*} = F(\underline{t}) = F(\underline{\bar{a}}+\delta\underline{\bar{p}}+\underline{e}) \simeq F(\underline{\bar{a}}) + L(\delta\underline{\bar{p}}+\underline{e}) = F(\underline{\bar{a}}) + L(\delta\underline{\bar{p}}) + L(\underline{e}) \quad (3)$$

Since we will be manipulating the perception $F(\underline{x})$ produced by a stimulus $\underline{x}$ rather frequently, I will denote the perception created by the stimulus $\underline{x}$ with a superscript asterisk, i.e. $F(\underline{x}) = \underline{x}^{*}$.

For ease of presentation I will discuss ~~"forced choice"~~ experiments in which either a test pattern, t, or an adapting pattern, a, is presented on each trial. Which of the two *is* presented on a given trial is determined randomly. The subjects' task in the experiment is to decide on each trial which of the two possible stimuli actually occurred. In other words, a stimulus x is presented where x is either a or t. The subject has to guess whether his perception, x^*, arose from presentation of t or a.

We will first consider the case in which the subject has complete knowledge of the appearance of t and a. We can assume, for example, that he gains knowledge of the relative appearances of t and a through feedback as to the correctness of his decisions during the course of the experiment or by being acquainted with a by itself and a suprathreshold exemplar of t.

However it is accomplished, we assume the subject has exact knowledge of the manner or direction in which $\bar{t}$ deviates from $\bar{a}$, and a knowledge of the "shape" of the cloud of deviations from average sensation created by each stimulus. As stated earlier, we will assume that the errors around $\bar{t}$ have the same distribution as the errors around $\bar{a}$. The local linearity condition, equation (1) above, implies that t^* is distributed with the same distribution, $L(e)$ around $\bar{t}^*$ as a^* is around $\bar{a}^*$.

A subject will perform best in such a task if on each trial he guesses that the x^* arose from that distribution which makes its occurrence most likely (9). This is the behavior assumed by standard PSDT. The only difference that I wish to emphasize is the many dimensional character of the situation. Figure 1 shows a three-dimensional representation of the subject's problem. The clouds around the points $\overline{a}^*$ and $\overline{t}^*$ represent the distribution of sensations arising from the presentations of a and t respectively. x^* represents the perception arising on some given trial. The subject's task is to guess which cloud the perception is really an element of. His best bet is to guess that a was presented if the probability of seeing x^* given that a was presented is higher than the probability of having the perception x^* if the stimulus was actually t.

The standard measure of discriminability in PSDT (labeled d') is the distance between the means of the distributions of the test and adapting stimuli divided by the standard deviation of the distributions. In the multidimensional case we have been discussing, it can be shown that

$$(d')^2 = (\overline{t}^*-\overline{a}^*)^T C^{*-1}(\overline{t}^*-\overline{a}^*) = (\delta\overline{p}^*)^T C^{*-1}\delta\overline{p}^* = \delta^2\overline{p}^{*T} C^{*-1}\overline{p}^* \qquad (4)$$

where C^{*-1} is the inverse of the subjective covariance function, and superscript T stands for transpose. This is the equation for a multi-dimensional ellipsoid(16). If the system is smooth in the sense discussed above, and the subject has good knowledge of the stimuli he is discriminating, contours of equal discriminability for compounds of two stimuli must lie on an ellipse. Such an experiment where

$$p = ap_1 + bp_2, \qquad a^2 + b^2 = 1$$

or more generally

$$p = \sum_{i=1}^{n} a_i p_i, \qquad \sum_{i=1}^{n} a_i^2 = 1$$

is called a subthreshold summation experiment.

That this performance can be achieved with human subjects is evident in the data of Figure 2.(17) These ellipsoids have been collected with widely differing stimuli by people who had no theoretical justification for the occurrence of ellipses. In two cases, they expressed surprise at the apparent need for a "square-law" or "energy" detector. In fact, no such operation is implied by this performance. The contour specified in (3) is unchanged if we take the square root of the equation, thus expressing the relationship in terms of Euclidean (more properly, Reimannian) distance rather than squared distance. The underlying "cause" of the observed behavior is simply the nature of the underlying statistical situation.

<u>The necessary existence of noise internal to the nervous system.</u>

Return with me now to equation (1). Were the equation true as it stands, we could learn nothing about the structure of the visual manifold in the neighborhood of the adapting stimulus other than its dimensionality. To see any other effects of its structure, a loss of information must occur in the system. In other words, the system must introduce noise which changes the statistics of the detection situation. The difference between the performance possible given the intrinsic noise in the stimulus and the actual performance possible by observers does shed some light on the processing taking place inside albeit rather dim and diffuse.

We can express the occurrence of this internal noise by replacing equation (1) with

$$F(\underline{t}) = F(\underline{a}) + L(\delta \underline{p}) + L(\underline{e}_e) + \underline{e}_i \qquad (5)$$

where $\underline{e}_e$ is the external noise associated with the stimulus and which can be measured by physical instruments, and $\underline{e}_i$ is noise from all sources internal to the system. The two noise sources, $\underline{e}_e$ and $\underline{e}_i$ are assumed distributed, with zero mean and covariance functions C_e and C_i respectively. The covariance function of a linear transformation of a random variable, in this instance $L(\underline{e}_e)$, is given by LC_eL^T. So the covariance function determining detection of perturbations around $\underline{a}$ is

$$C^* = LC_eL^T + C_i \tag{6}$$

Since there are only a finite number of components in the nervous system and in particular only about a million fibers in each optic nerve, the range of F is finite dimensional and $\underline{e}_i$ is finite dimensional. Therefore, C_i can be represented by a matrix. It is apparent that if C_i is identically zero, which means that $\underline{e}_i$ is invariably just a vector of zeroes, then eq. (5) reduces to eq. (3) and (6) reduces to

$$C^* = LC_eL^T \tag{7}$$

Let us consider first the contention that were there no internal noise, i.e. that $C_i = 0$, nothing could be learned about the properties of F in the neighborhood of $\underline{a}$ other than its dimensionality. For the sake of concreteness, consider the transformation which takes the visual patterns to the outputs of retinal ganglion cell outputs.

The pattern in the outside world is an infinite dimensional vector as mentioned earlier - an infinite number of values, one for each point in a region of the plane - are needed to specify a particular pattern. The pattern of outputs of the ganglion cells is given by $F_{eye} = T_g \circ T_r \circ T_o$ which were defined earlier. The reason for using this example is that the greatest loss of information in the system occurs between the outside and the optic nerve. In particular, all visual information must go through the bottleneck of the optic nerve. No amount of elaboration after that point can restore the lost degrees of freedom.

The output of this transformation is the firing rates of the million or so axons in the optic nerve. Suppose we could attach each of these optic nerve fibers to a viewing screen so that the brightness of each point on the screen was determined by the firing rate of one fiber. We would see a pattern $\underline{y} = F_{eye}(\underline{x})$ for any pattern of light, $\underline{x}_o$ cast on the retina. Assume F_{eye} as just described is totally deterministic.

It is obvious that F_{eye} is singular; it takes patterns from an infinite dimensional space and collapses them into a space of around 10^6 dimensions. Therefore, there must be an infinite number of patterns which are indistinguishable simply because they get mapped to the same output pattern - they are not resolved by the system. We can separate the patterns which are resolved from those which are not, in the following way.

We need to work first with non-stocastic patterns before investigating what happens to ensembles of patterns. Given a non-stochastic pattern, $\underline{\bar{a}}$, we will call the set of all patterns $\underline{\bar{x}}$ for which $F_{eye}(\underline{\bar{a}}+\underline{\bar{x}}) = F_{eye}(\underline{\bar{a}})$ (equivalently $L_{eye}(\underline{\bar{x}}) = 0$) to be the null space or kernel of F_{eye} at $\underline{\bar{a}}$ or equivalently the null space of

$L_{eye}(\bar{\underline{x}})$. We will label this space $N_{\bar{\underline{a}}}$. We will call the orthogonal complement of the null space $C_{\bar{\underline{a}}}$, the co-kernel of F at $\bar{\underline{a}}$. In this way we have divided the total space of stimulus patterns into two unique subspaces such that any pattern $\bar{\underline{x}}$ can be expressed as the sum of a component, $\bar{\underline{x}}_n$ from the null space and $\bar{\underline{x}}_c$ from the co-kernel, $\bar{\underline{x}} = \bar{\underline{x}}_n + \bar{\underline{x}}_c$. $F_{\bar{\underline{a}}}$ is an isomorphism when restricted to $C_{\bar{\underline{a}}}$. Patterns of the form $\bar{\underline{s}} = \bar{\underline{a}} + \bar{\underline{x}}_n$ which contain no component from $C_{\bar{\underline{a}}}$ obviously cannot be resolved from $\bar{\underline{a}}$. Their thresholds are arbitrarily high, noise considerations aside. Conversely, were it not for noise, any pattern in $C_{\bar{\underline{a}}}$ should be discriminable from any other.

However, there is noise in the stimulus. Equation (3) represents this situation. We only have to concern ourselves with the portions of the patterns in $C_{\bar{\underline{a}}}$, so from here on we will tacitly assume that we are ignoring the components of patterns in $N_{\bar{\underline{a}}}$. According to eq. (3) we can write

$$\underline{t}^* - \underline{a}^* = (\delta\bar{\underline{p}})^* + \underline{e}^* = L(\delta\bar{\underline{p}}) + L(\underline{e})$$

so $(\delta\bar{\underline{p}})^*$ is centered at $L(\delta\bar{\underline{p}}) = \delta L\bar{\underline{p}}$. Combining eq. (7) and eq. (4) we have

$$\begin{aligned}(d')^2 &= \delta^2 (L\underline{p})^T (LC_e L^T)^{-1} L\underline{p} \\ &= \delta^2 \underline{p}^T L^T L^{T^{-1}} C_e^{-1} L^{-1} L\underline{p} \\ &= \delta^2 \underline{p}^T C_e^{-1} \underline{p} \qquad (8)\end{aligned}$$

giving us the result we were seeking. If there were no noise internal to the system, the threshold loci would be determined by the statistics of the stimulus. All that could be distinguished would be the dimensionality of the kernel and the co-kernel. That this dimensionality can be determined without reference to the statistics of the system is illustrated by the many ways it can be shown that our color perception is three dimensional.

Of course it is pointless to try to determine the dimensionality of a system as complex as the visual or auditory systems, however, it is important to note that their dimensionality can be no higher than the number of fibers leading from the receptors to the brain. In certain systems such as smell, where little is known even about the varieties of receptors extant, psychophysical measures could perhaps determine just how many there are.

The assumption of a noiseless system is nonsense, of course. Even the viewing screen analogy drawn earlier is unrealistic in disregarding the noise of the light coming from the screen. We will now examine the effect of internal noise on the performance of the system and show how, in the visual system its presence gives rise to the modulation transfer function.

As indicated earlier, equations (5) and (6) express the distributions governing detection when the system contains internal noise. Figure 3 illustrates, in three dimensions, the relationships we are considering. Several simplifications have been made to clarify the picture. First, the distributions of $\underline{e}_e$ and $\underline{e}_i$ are assumed to be spherical. That is, their covariance functions are assumed to be of the forms $\sigma_e^2 I$ and $\sigma_i^2 I$ where I is the identity function and σ_e^2 and σ_i^2 are the external and internal variances respectively. Secondly, the locations of the means of two test patterns are represented, but their distributional clouds (which are assumed identical to $\underline{a}$'s) are elided. An external adapting stimulus distribution and the means of two arbitrary test patterns are represented in 3a. These patterns are transformed by F to the pattern in Figure 3b. This picture illustrates the situation if there were no noise internal to the system. $\bar{\underline{t}}_1$ has been transformed to $\bar{\underline{t}}^*_1$ which is further away from $\bar{\underline{a}}^*$ than $\bar{\underline{t}}_1$was from $\bar{\underline{a}}$, and conversely $\bar{\underline{t}}^*_2$ has been squished closer to $\bar{\underline{a}}^*$. However, the standard deviation of noise in the direction of $\bar{\underline{t}}^*_1$ has increased, and in the direction of $\bar{\underline{t}}^*_2$ decreased in proportion to their own movement so the statistics determining their detection remain unchanged. Now, if the internal noise represented in 3C adds to the transformed external noise, the governing distribution is as indicated in 3d. $\bar{\underline{t}}^*_2$ has been virtually buried in the noise, whereas $\bar{\underline{t}}^*_1$ is still relatively discriminable.

Notice that as the ratio of the external noise variance to the internal decreases, the equation controlling detection ranges from equation (7) to

$$
\begin{aligned}
(d')^2 &= \delta^2 \underline{p}^{*T} C_i^{-1} \underline{p}^* \\
&= \delta^2 \underline{p}^T L^T C_i^{-1} L \underline{p} \qquad (8)
\end{aligned}
$$

In this case, which holds when the external noise is negligible, threshold contours are determined by the matrix $L^T C_i L$. Even here there is ambiguity between the shape of the internal noise distribution and the transformation the stimulus undergoes. To make assertions about L, we must make assumptions about C_i. It can be demonstrated that if we replace L with $G^{-1}L$ and C_i with $G^{-1}C_i G^{T^{-1}}$ the quadratic form determing thresholds remains constant. This is true whatever the ratio of external and internal noise. We might as well assume the canon that $C_i = \sigma_i^2 I$, in other words that the distribution of internal noise is spherical and all differences in the discriminability of stimuli are due to the transformation from the world to percept. ~~In physiological terms, the assumption means that the noise in each neuron's firing ratio is about the same and the noise components are uncorrelated from neuron to neuron.~~ In any case, it is important to recognize that since the "true" state of events in the external environment is not available internally, there is no relevant basis against which to assert a particular form for C_i. All that can matter is the relation between L and C_i.

An illustration of the points made in this section is provided by the change in shape which the contrast sensitivity function (CSF) undergoes with change in luminance level. A CSF is the graph of the reciprocal of the contrast required for detection of a grating versus the spatial frequency of the grating. In general, the luminance waveform of the grating is a sinusoid, as is the current case. However, later we will be comparing CSFs obtained with square waves to those obtained with sines. Since there is not the space to go into detail here, I will just point out that the loss of distinctive shape, (particularly the loss of the low frequency roll-off which is necessarily of neural origin) in the lower luminance curves in Figure 4 can be attributed simply to the change in the relative amount of neural noise and external quantal noise which occurs with charges in light flux. Specifically, if a compressive transformation such as the logarithmic takes place at the receptors, the variance of noise which is pointwise independent will decrease relative to neural noise which occurs after the signals are mixed by neural processing. If the receptor transformation is logarithmic, and if we assume that neural noise central to receptive field transformations etc. is relatively constant (18), the change from the de Vries-Rose law to Weber's law is explained at the same time as the change of shape of the CSF (19).

From this point on we will consider only visual experiments conducted under luminance levels at which external noise is insignificant.

The measurement of local eigenfunctions

We now return to the question of the viability of the notion of feature detectors as a psychophysical explanatory device. Assuming that external noise is insignificant, and the subject has full knowledge of the stimuli, we have seen that

$$(d')^2 = \delta^2 \underline{p}^{*T} C_i^{-1} \underline{p}^* = \delta^2 \underline{p}^T L^T C_i^{-1} L \underline{p} \tag{9}$$

I now wish to point out several properties of this quadratic form. It is apparent that the information about the system we can gain by varying $\underline{p}$ depends only on the total lump $L^T C_i L$, not the individual functions producing it. I have already indicated how C_i can be replaced with any other covariance matrix if a complimentary change is made in L. In particular, there exists a matrix K such that $C_i = KIK^T$ and we can express our threshold form as

$$(d')^2 = \delta^2 \underline{p}^T L^T (KIK^T)^{-1} L \underline{p} = \delta^2 \underline{p}^T Q^T I Q \underline{p} = \delta^2 \underline{p}^T Q^T Q \underline{p} \tag{10}$$

where $Q = K^{-1}L$.

For simplicity we will work with equation 10 and consider Q to be the matrix specifying the overall action of F near our adapting stimulus. To tie what we have done back to the original issue of the viability of feature detectors, we can interpret Q to be a matrix whose rows are some set of proposed detectors. Figure 5 illustrates such a matrix pictorially. Each column corresponds to a location on the retina. Each row is one putative detector. As examples, the response functions of "bar-like", and "edge-like" and "Fourier" detectors are illustrated. The output of each detector is given by the inner product of the response function of the cell and the pattern on the retina. Assume there are n detectors which are combinations of m receptors so Q is n x m. Notice first that since Q^TQ is m x m we can never tell that the number of detectors is greater than the number of receptors, reillustrating the point made more generally earlier that the dimension of the system cannot be shown to be greater than that of the tightest information bottleneck in the system.

Next note that for any orthonormal matrix O,

$$Q^TO^TOQ = Q^TIQ = Q^TQ$$

due to the defining quality of orthonomal matricies -- their transpose is their inverse. So, for any proposed set of detectors, Q, there is an infinite class of sets producing identical threshold data. This class is called the similarity class of the matrix Q.

So, the most we can determine on the basis of threshold experiments is this similarity class. I mentioned above that a quadratic form of the sort we are discussing defines a multi-dimensional ellipsoid. The similarity class of Q is simply the class of all systems which produce exactly the same ellipsoid as Q does. Of the continuum of forms which could generate the observable data, it is useful to choose one which presents the information in an intuitive manner and which allows comparison between similarity classes.

A very appealing description is simply to specify the orientation and length of each of the axes of the threshold elipsoid. The principal axis theorem (16,20) which is the underpinning of factor analysis tells us that any symetric matrix such as $S = Q^TQ$ can be written in the form E^TDE where E is an orthonormal matrix whose rows are the eigenvectors of S and D is a matrix of whose main diagonal contains the corresponding eigenvalues and contains zeroes everywhere else. These eigenvectors are the axes of the elipsoid, each of whose length is given by twice the square root of its associated eigenvalue. These sets of eigenvalues and vectors summarize all that can be learned about the steady state performance of a sensory system near some particular adapting stimulus.

The experiment of Campbell and Robson

We now have the tools to consider the experiment by Campbell and Robson which was the original motivating influence behind this article. The rationale behind the use of trigonometric (Fourier) functions for the analysis of systems can now be placed in its proper context.

There is a very common class of linear systems, which we will call invariant linear systems, for which sines and cosines are the eigenfunctions. In our terms, these systems are those whose similarity class contains a matrix whose rows are identical to each other except for a shift in position. For instance, if the matrix illustrated in figure 5 contained only "bar" detectors like those shown in the first and second rows and all were identical to each other except for the retinal position upon which they were centered, then sines and cosines of various frequencies would be the eignevectors of the matrix, and the CSF which gave the eigenvalues associated with each frequency would contain all the information needed for complete description of the performance of the system.

In the experiment upon which they based most of their theorizing they determined subject's thresholds for perceiving bar patterns on a television-like screen. The luminance profile of the bar pattern was either a sinusoid or a square wave. The frequency of the bars ranged from about 0.2 to 40 cycles per degree visual angle. The subject's task was to adjust the contrast of the grating pattern until it was barely discriminable from a totally blank screen. In other words, they collected CSFs for both sinusoidal and square wave patterns. In our notation the adapting pattern, a, was a blank screen of some fixed luminance, ℓ. The test pattern was either of the form

$$\underline{t}_c(\omega) = \ell + \delta \cos(\omega x) \tag{11}$$

or

$$\underline{t}_s(\omega) = \ell + \delta \, \mathrm{sgn}(\cos(\omega x)) \tag{12}$$

where sgn(x) = +1 if $x \geq 0$, −1 otherwise, $\underline{t}_s$ can also be expressed in terms of sinusoids by the infinite Fourier series

$$\underline{t}_s(\omega) = \ell + \delta \frac{4}{\pi} \sum_{n=1,3,5,\ldots}^{\infty} \frac{1}{n} \cos(n\omega x) \tag{13}$$

in this latter form $\underline{t}_s$ can be considered a stimulus in a subthreshold summation experiment as defined earlier.

Figure 6 shows their data and some simulations of it I have added. The open circles in the upper portion of the figure are their obtained CSF for cosine gratings. The solid line through these points is a least squares fourth degree polynomial fitted to their data for purposes of testing the notions at issue here. The other curves I have added to this figure are calculated with no free parameters from this estimate and any errors in it are propagated into the derived curves. The filled squares in the upper portion of the picture form their CSF for square wave gratings.

In the lower portion of the Figure the ratios of the sensitivities for the square waves to sine waves are plotted with filled circles. Line segments indicate their standard errors around these ratios. They have placed a solid horizontal line through $4/\pi$ which is the ratio of the fundamental cosine component of $\underline{t}_s(\omega)$ to $\underline{t}_c(\omega)$. The final line of their original illustration is the dashed line which rises to the left above their obtained ratios. This line gives the ratios of the amplitudes of the stimuli which would be predicted on the assumption that "the threshold is determined by the peak value of the function obtained by passing the modulation wave form of the grating through a filter whose attenuation characteristic has the form of the contrast-sensitivity function" (3,p.563).

The ratio they compute under this model for each frequency ω is

$$\frac{\max\limits_{x} \frac{4}{\pi} \sum\limits_{n=1,3,5..} \frac{1}{n} \cos\,(n\omega x)\,CSF(n\omega)}{CSF(\omega)}$$

They propose this model as a straw man against which they contrast their data. The notion they proposed in its stead is "that the visual system behaves not as a single broad-band spatial filter but as a number of independent detector mechanisms each preceded by a relatively narrow-band filter 'tuned' to a different frequency" (3, p564). No specific predictions can be made from this model, and I won't pursue its several logical problems any further here.

The last two curves in the figure are predictions based on the ideas developed in this paper. They rest on two assumptions: first, and most solidly, the general theory of threshold detection under conditions of stimulus certainty presented above; secondly, The assumption that sinusoids a factor of 3 apart in frequency would be orthogonal with respect to the system. Two patterns, $\underline{p}_1$ and $\underline{p}_2$ are orthogonal to one another if $\underline{p}_1^T\underline{p}_2 = 0$. Two patterns are orthogonal with respect to the system G if $\underline{p}_1^T G^T G \underline{p}_2 = 0$. The relevance of this notion in the present situation is that the detectability of the sum of two patterns is determined by

$$(\underline{p}_1+\underline{p}_2)^T G^T G(\underline{p}_1+\underline{p}_2) = \underline{p}_1^T G^T G\underline{p}_1 + \underline{p}_2^T G^T G\underline{p}_2 + 2\underline{p}_1^T G^T G\underline{p}_2 \qquad (14)$$

This expression can be readily extended to combinations of more than 2 patterns such as the infinite sum representing the square wave in our current example. Only if the pattern components are orthogonal with respect to the system can the square of the detectabiltiies of two patterns be added to yield the square of the detectability of the sum.

These assumptions yield thresholds for the square waves given by

$$\sqrt{\frac{4}{\pi}\sum_{n=1,3,5,\ldots}\left(\frac{4}{n\pi}\,CSF(n\omega)\right)^2}$$

The line through the square wave CSF is the curve predicted using this expression and the line through the ratios is the corresponding predictions for them. The apparent better fit for the ratios is probably due to the cancelling of some of the error in the polynomial fit of the raw data. <u>All</u> harmonics for which there is any appreciable response by the visual system contributed to the detectability of the pattern. For the lowest frequency, even the addition of the 21st harmonic to the calculation made a slight improvement in the fit. This is not surprising since the perceived pattern <u>is</u> the entire sum, not some arbitrary portion of it.

Note that the validity of the second assumption made above is fortuitous. Other data indicates adjacent frequencies are not orthogonal vis a vis the system (21), and in fact frequencies even a factor of 2 apart are not generally found to be orthogonal. These dependencies have been interpreted as demonstrating the "tuning" of the various potential mechanisms. What they actually demonstrate will be discussed elsewhere.

It should be noted that Reimannian metrics of the form given here are the only metrics in which the norm of a pattern is not dependent upon its representation. In other words, we arrive at identical detectabilities for square waves whether we represent t_s by either expression 12 or 13. Such systems are the only context in which Fourier analysis makes sense, so all those experiments which have shown the applicability of Fourier techniques to psychophysical phenomena lend support to the thesis presented here.

Effects of uncertainty

I am inclined to lump all ways in which data deviates from the basic theory presented here into just 3 categories: a)deviations of the perceptual manifold from linearity for larger signals, b)true non-differentiabilities and catastrophies on the manifold, and, c)the numerous effects of uncertainty mentioned below. The first of these is not a major conceptual problem. It is merely what requires us to use the notion of a manifold rather than a linear space. Simple nondifferentiabilties probably are obscured by the stochastic nature of things even if they do occur. However, discontinuities or catastrophies and the associated phenomenom of multiple steady states are quite common, particularly in binocular vision. I discuss the problem of distinguishing between them and effects I would rather label uncertainty elsewhere (22). Here I will only briefly discuss the third catch-all category - uncertainty.

Books need to be written delineating the multiplicity of effects engendered when either the stimulus situation, the subject's expectations, or the inprecision of his memory change the appropriate statistical model. Many of the stimulus paradigms which have been used in the past have corresponding statistical models which are nearly intractable. Each such situation requires detailed analysis before any solid predicitons can be made. The general signal detection equations for these problems are given in Middleton (8) and I'm sure a large number of specific situations have been examined by workers in communication theory in the intervening years.

Rather than developing any particular examples, I will just state a few general conclusions.

1) uncertainty effects are ubiquitous. (4,23,24) list a few papers demonstrating them.

2) The slopes of the psychometric functions obtained generally increase relative to their means as uncertainty increases (23). The simplest case to handle is that in which uncertainty is subspacewise - that is, the pattern to be detected is specified only down to a subspace. For example, if the pattern to be detected is simply of the form $\underline{p} = M\,\underline{a}$ where M is a matrix whose columns are a set of orthogonal, unit length patterns and $\underline{a}$ is a random vector of coefficients for the component patterns. In this case $(d')^2$ will be distributed as a χ^2_n. It is well known that the standard deviation of these distributions increases as just $\sqrt{n}$ while the mean equals n.

3) Contrary to popular dogma, it is usually a bad idea to randomize the conditions of an experiment. Randomization causes effects due to increased uncertainty (24) and differential amounts of uncertainty between conditions containing stimuli of differing degrees of complexity (4).

Critical Bands and similar masking phenomina

The final topics I will discuss are the phenomenon of critical bands and a related observation by Campbell & Robson. I will discuss Campbell & Robson's effect first, then build upon its explanation to explain in the discrimination of a pattern from a background of randomly selected members of any specified set.

Campbell & Robson examined "...the contrast level at which... square-wave gratings can be distinguished from sine-wave gratings having the same spatial frequency and the same fundamental amplitude. A display was devised which alternated at quarter second intervals from a sine-wave to a square-wave. "The experimenter set the sine /square ratio to $4/\pi$ and the subject was provided with a control which enabled him to adjust the contrast of both gratings together while maintaining the ratio of their contrasts constant. At each pre-set spatial frequency the subject raised the contrast of both gratings together until he could distinguish one grating from the other" (3,560-561).

Figure 7 shows the geometry of the situation. It may also help clarify some of the concepts presented earlier. Frequencies a factor of 3 apart are apparently orthogonal. Thus it is appropriate to plot the fundamental vector perpendicularly to that representing all the other components. Campbell & Robson present data only for frequencies above approximately 1 cycle per degree visual angle. It is apparent from Fig. 6 that at these frequencies, harmonics above the third are relatively insignificant. Fig. 7 represents the 3 cycle per degree case. Even at this frequency the sensitivity for the third harmonic is only about 1/3 that for the fundimental. So as can be seen in the figure, it appears that the square wave comes to threshold about when the fundimental alone does. In their discrimination experiment, Campbell & Robson increase the contrast of the square wave (increasing amplitude along the diagonal line), while constantly comparing this stimulus against an adapting pattern consisting of the fundamental components. From the picture it can be seen that one would expect the two patterns to be distinguishable only when the portion of the square wave orthogonal to the fundamental reaches its own threshold. This is what they found. Actually, the most interesting aspects of their confirmation of this picture are the facts that it shows that there is nothing sacred about the blank field adapting pattern, and that the system is linear within experimental precision in a considerable region in pattern space around the blank pattern. The general principle that two patterns become discriminable when their vectorial difference reaches threshold explains a great variety of masking phenomena in which alteration of the system due to adaptation is minimal.

We finally turn to the explanation of critical bands. If a pattern is presented against a background of noise it becomes more difficult to detect. Also, as indicated earlier, its detectability becomes more dependent upon the statistics of this external noise. However, what counts in detection is the variance of the noise "in the direction of" the pattern to be detected (cf. figure 1). In audition, it was observed a pure tone was harder to detect when presented in white noise. However, when a sufficiently wide gap in the spectrum of the noise is made around the frequency of the tone, the detection of the tone is virtually unaffected by the presence of the noise. This gap is called the critical band (25). Analogous effects have been found in vision (26).

The principle invoked in explanation of these effects is straightforward -- the distribution of the noise with a gap in it forms a cloud in a hyperplane in pattern space. If the pattern to be detected is within this hyperplane, its detection is determined by the noise distribution. If the pattern is orthogonal to the hyperplane, its detection is unimpaired by the noise. If neither, its detectability is dependent mainly on the component of the pattern orthogonal to the plane of the noise. Thus, if the pattern is orthogonal to the noise,its perception is unimpaired.

Summary

I have presented here the outline of a general theory of threshold behavior and indicated several points of correspondence between it and available data. Hopefully, the overall picture (which seems to me to have almost the status of a tautology) will not be lost in debates over details which may require modification. This theory in no way denies the complexity of psychophysical phenomena. However, it does elaborate severe restrictions on the interpretations which can be legitimately based on these data. Nor is it meant to trivialize the insights which can be gained from the introspective technique of psychophysics. The present paper is itself a product of that tradition. Indeed, an implication of the ideas presented here is that rather sophisticated statistical models are required to describe a subject's behavior in even the simplest experiment.

Finally, a comment on experiments which contain a temporal component, eg. studies of forward and backward masking and motion perception. The analysis of these situations is considerably more complex. However, it is not fundamentally different in principle. These problems of discrimination can be pictured either as problems of discriminations of paths on the sensory manifold as described here or, more adequately, as the discrimination of points on a manifold in the cartesian product of the space of static patterns considered here and the time axis. In any event, the character of the data obtained and the implications which can be supported by them is unchanged.

FIGURE CAPTIONS

1. Representation in three dimensions of the subjects' dilemma in a threshold experiment. Each point represents the perception of the subject on a given trial. The cloud centered on $\vec{\underline{a}}^*$ consists of percepts arising from presentation of the adapting stimulus a : the elements of the cloud centered on $\vec{\underline{t}}^*$ are percepts arising from presentation of the test stimulus t . On a given trial, the subject's percept may, perchance, correspond to a point at $\underline{x}^*$. The subject must decide whether the experience is more likely to be a member of the distribution generated by a presentation of a or of t.

2. Some elipses found in various fields of psychophysics. (17).

a). Thresholds for luminance increments and decrements of small circular targets presented binocularly. The luminance change in the right eye is plotted horizontally; that in the left, vertically. Data from two subjects is shown.

b). MacAdam color discrimination elipses plotted on the standard cromaticity diagram. Elipses are plotted ten times actual scale.

c). Plots of the relationship between two brief flashes at threshold with different intervals between them. The interval between the flashes in msec. is marked near the origin of each curve. From Rashbass.

3. A 3-dimensional illustration of the effects of the transformations performed by and noise added by the visual system. See text for explanation.

4. Contrast sensitivity function for various levels of retinal illumination (measured in trolands). (19)

5. Representation of the product of a hypothetical matrix of "detectors," Q, and an arbitrary pattern, p, in the visual field. Only one dimensional profiles of the functions and patterns are represented. Each row represents the spatial response function of one detector. From the nasal, N, to

temporal, T, sides of the visual field. The response of each detector is given by the inner product of the detector function and p. The total output pattern is given by the vector of the outputs of all the detectors. The first two detectors represented are similar in form to proposed "bar detectors." They are identical to each other except for position with respect to the visual field. The third function is similar to "line detectors." (2) The last two functions are like proposed sine and cosine detectors.(5)

6. The contrast sensitivity functions obtained by Campbell and Robson (3) for sine (open circles) and square (filled circles) wave grating patterns. See text for discussion.

7. The geometry of Campbell & Robson's sine-square wave discrimination experiment. The picture is scaled to correspond to their three-period per degree data. a is the point representing the blank adapting pattern. f_t and h_t are the thresholds for discriminating the fundamental component (the 3c/d sinusoid) and the higher harmonics (mainly the 9c/d component) from a respectively. Square waves are produced (approximately) by the vectorial addition of 1/3 of the 9c/d component to the 3 c/d component. The diagonal line represents square waves of varying amplitude with threshold reached at S_t. To discriminate the square wave from its fundamental component, however, requires detecting the component of the square wave orthogonal to the fundamental component. This component is not discriminable (as indicated by the dashed vectors) until d_t when the threshold found for discriminating those components from a blank pattern is reached.

FIGURE 1

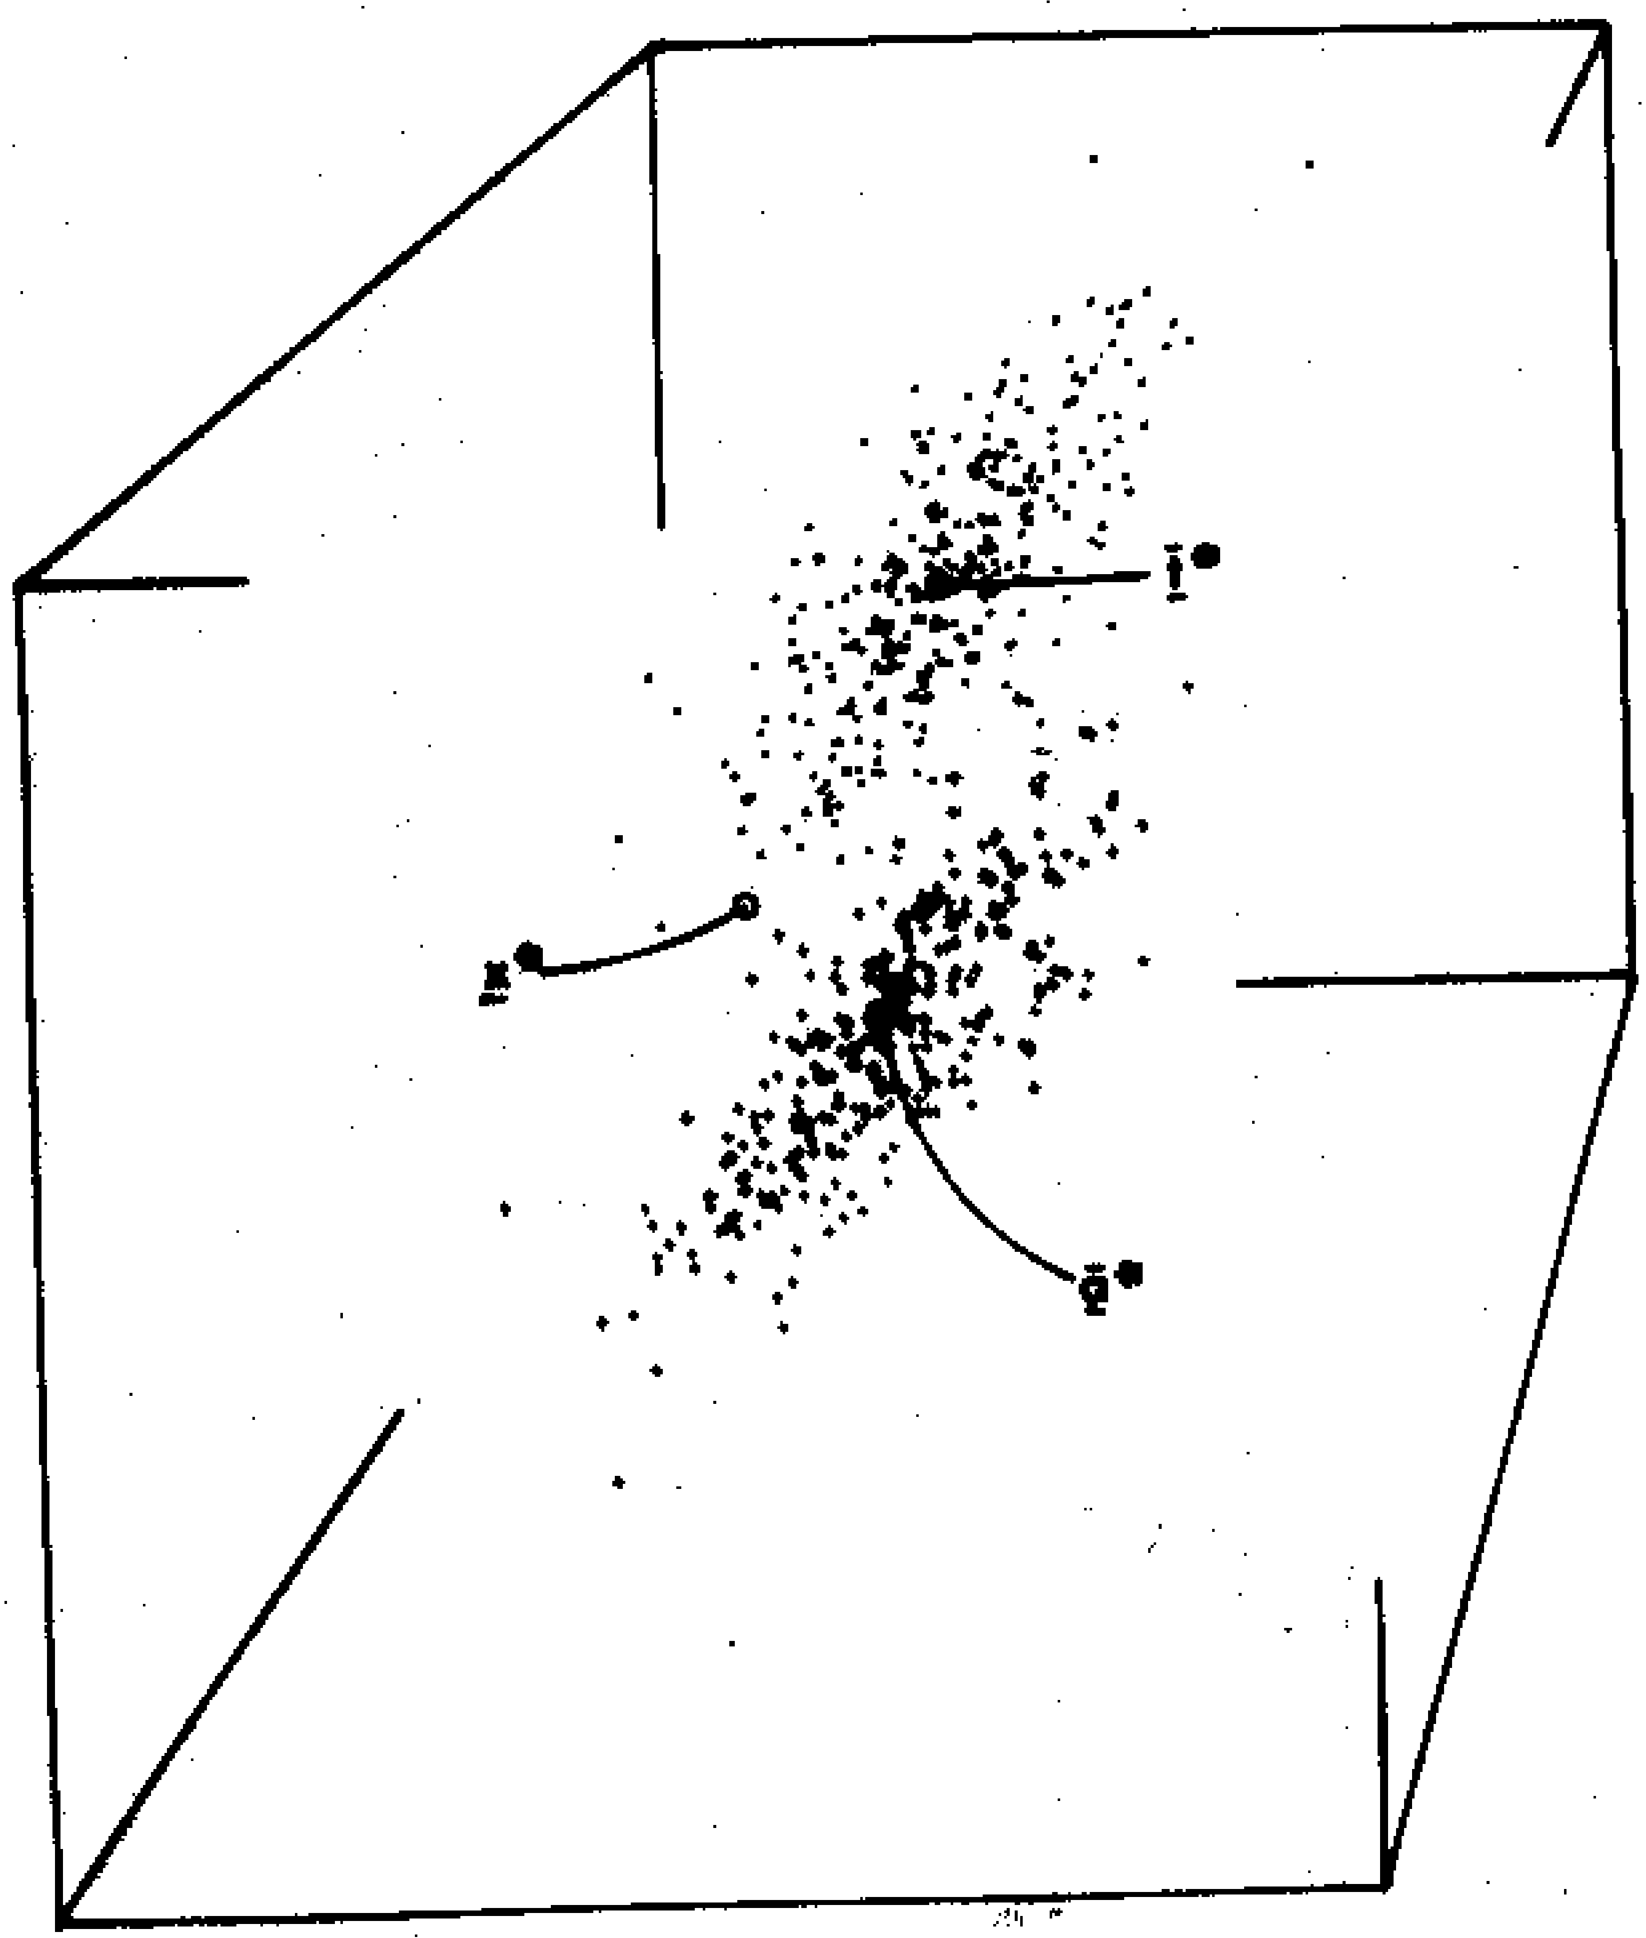

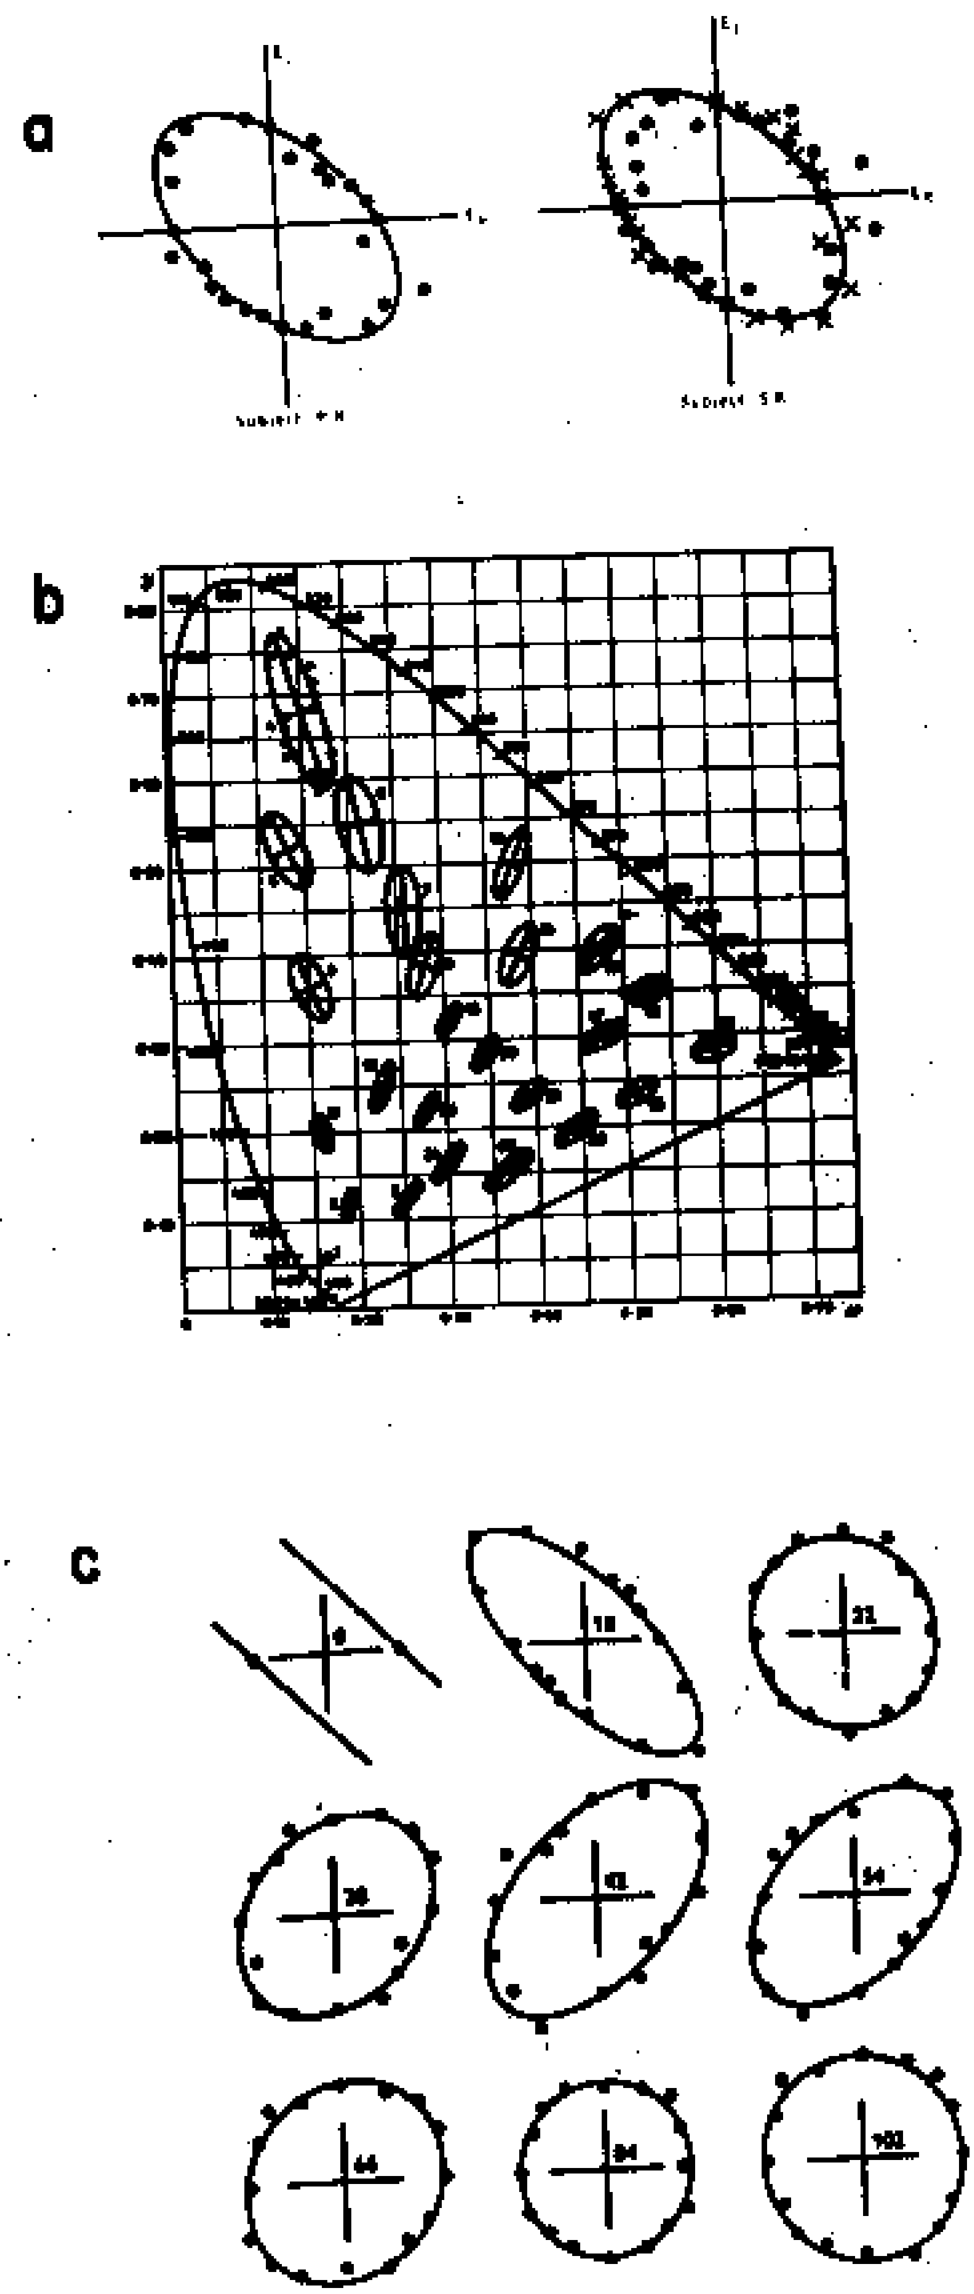

FIGURE 2

FIGURE 3

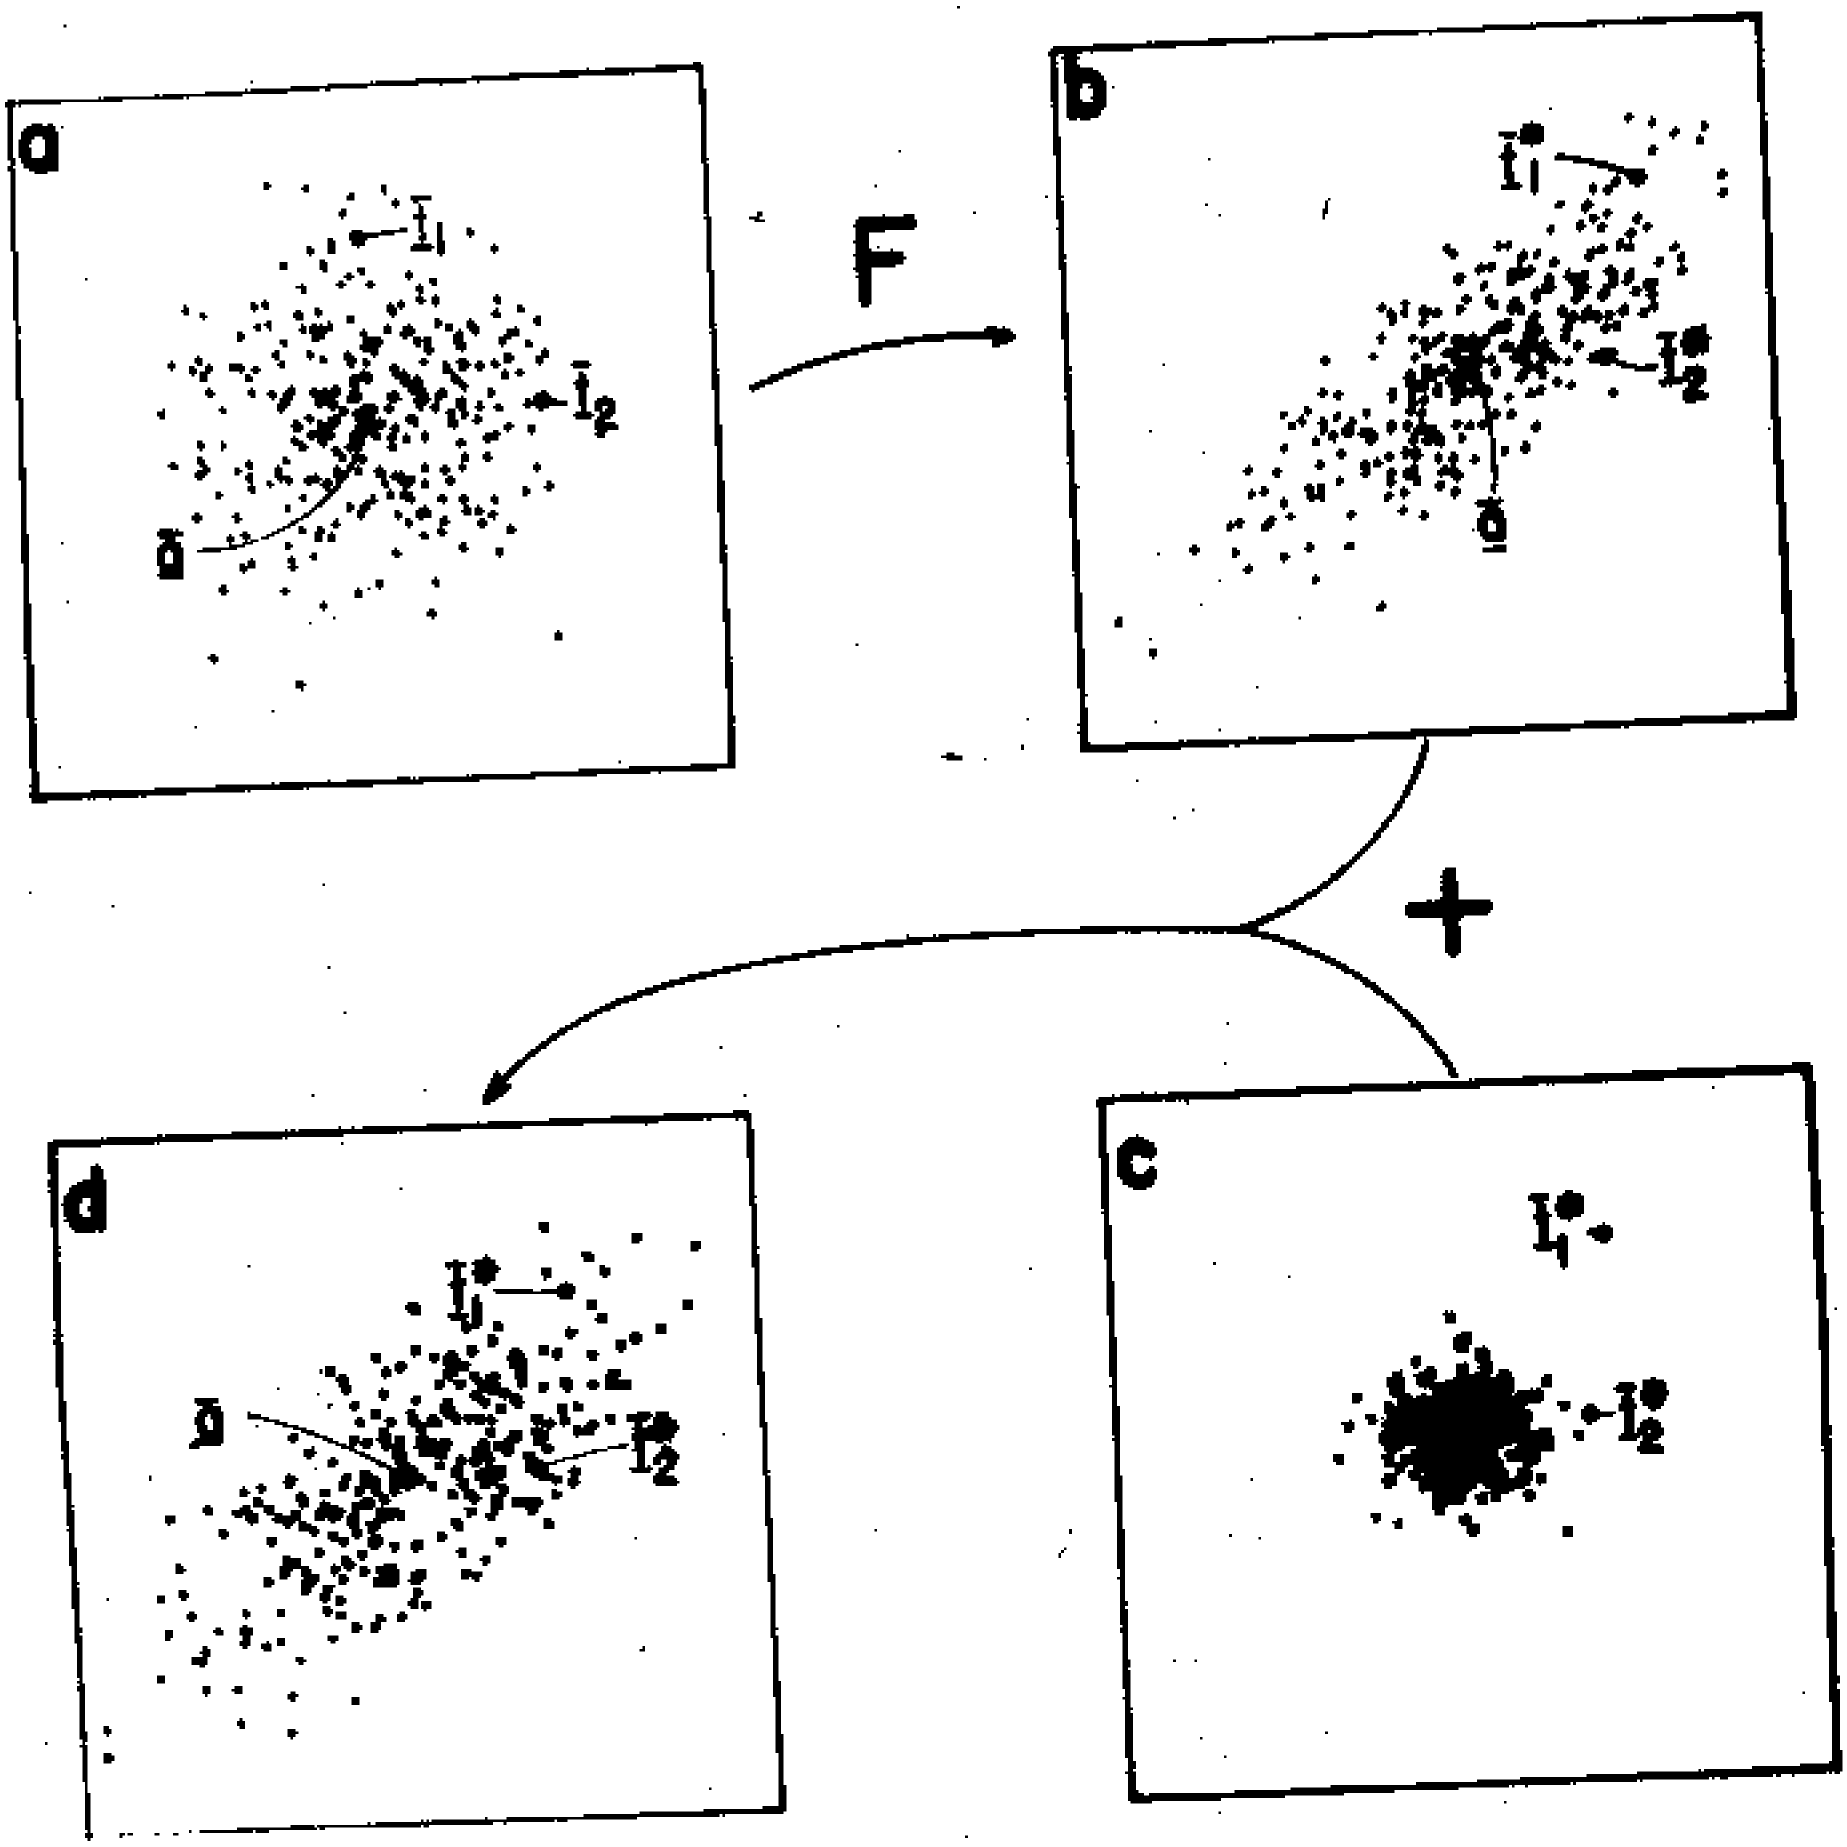

FIGURE 4

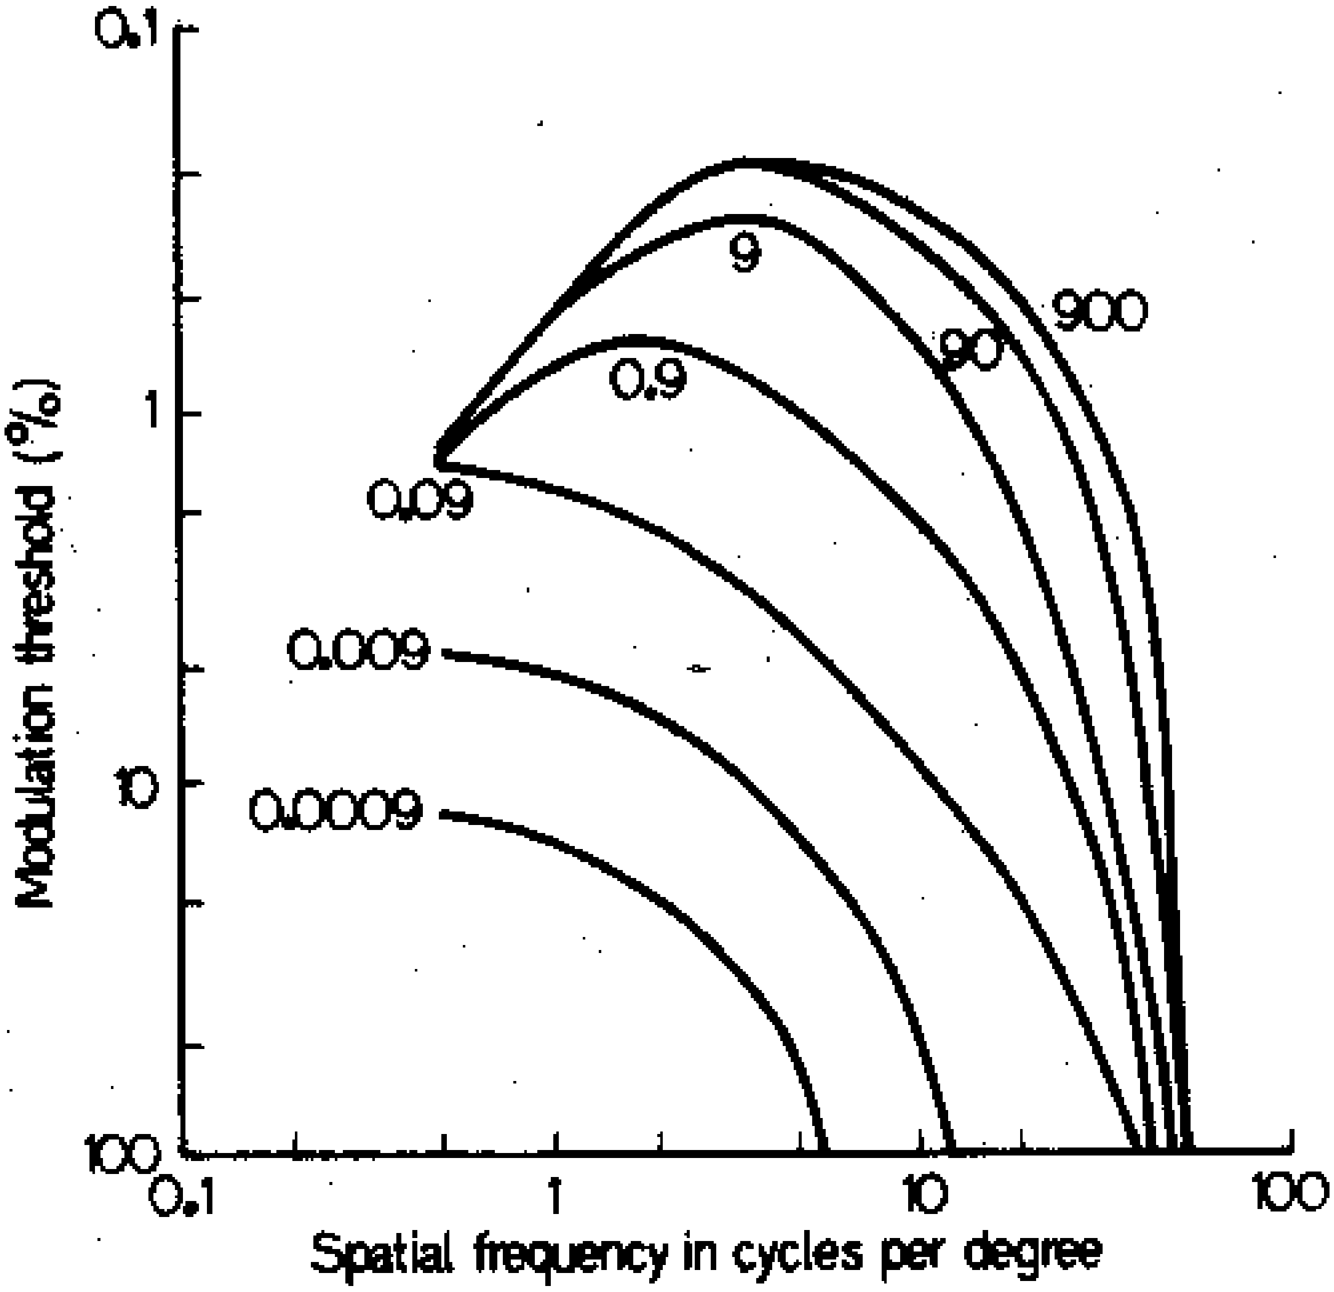

[redacted]

FIGURE 5

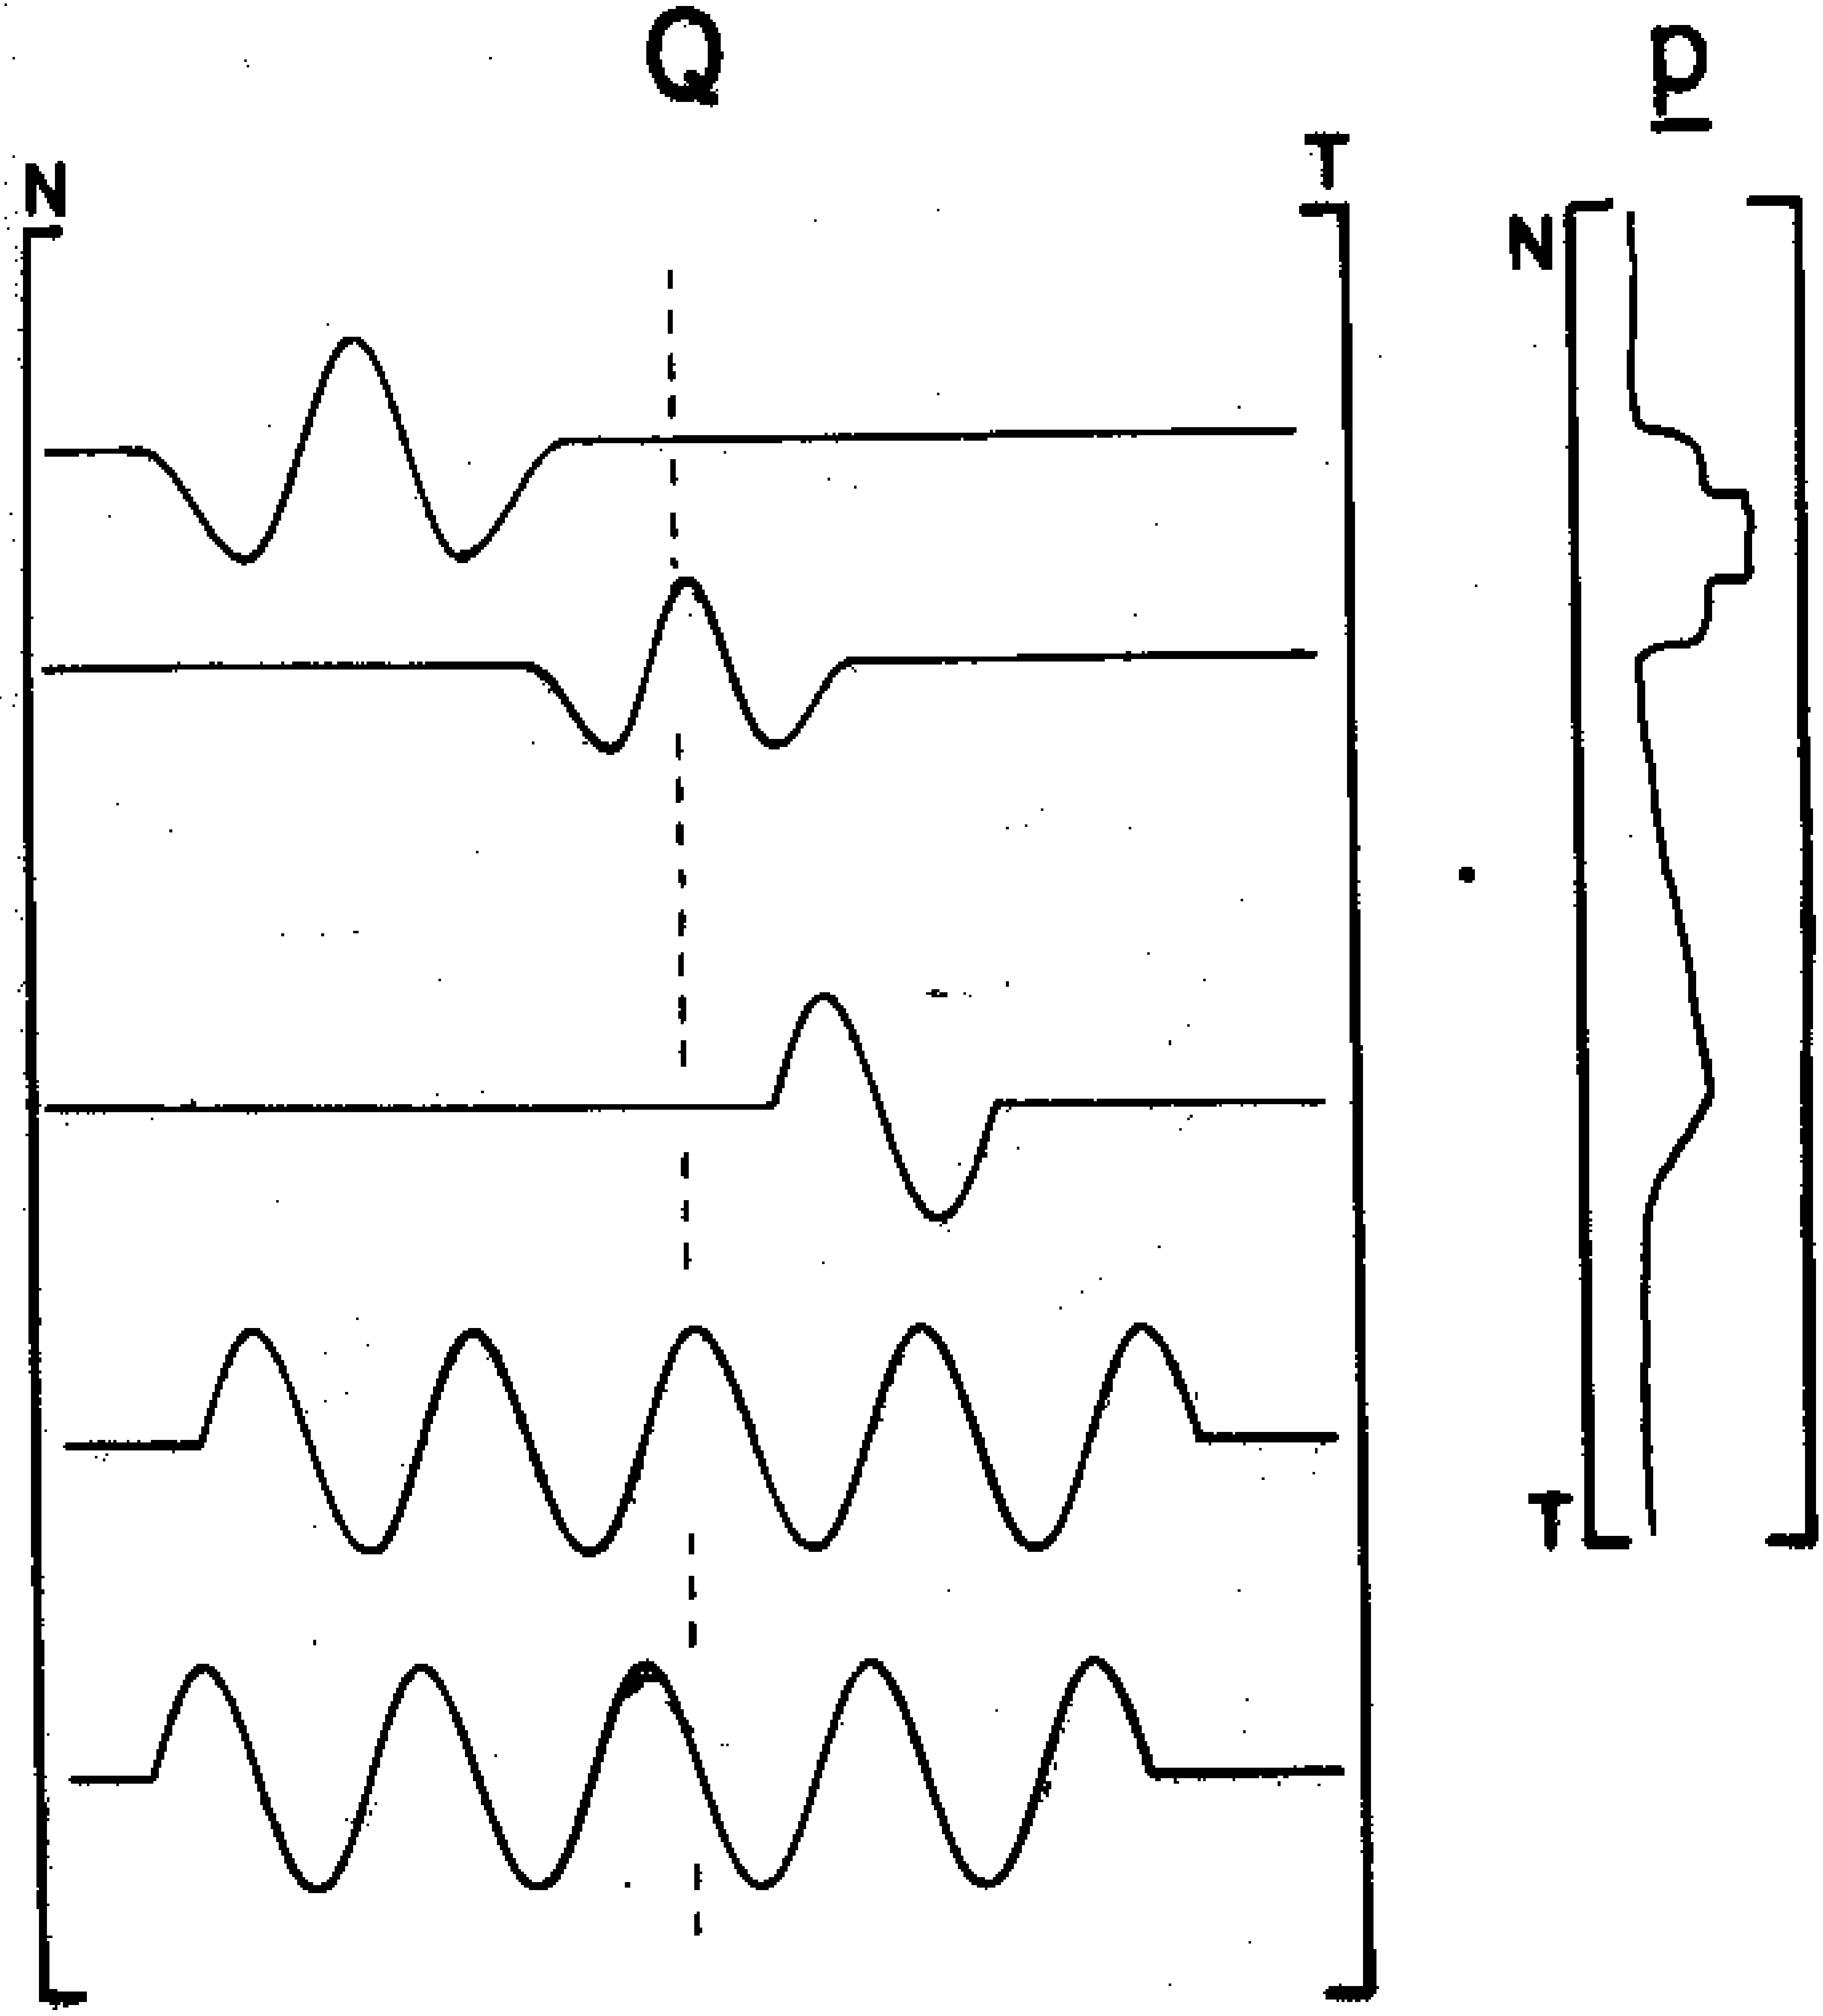

FIGURE 6

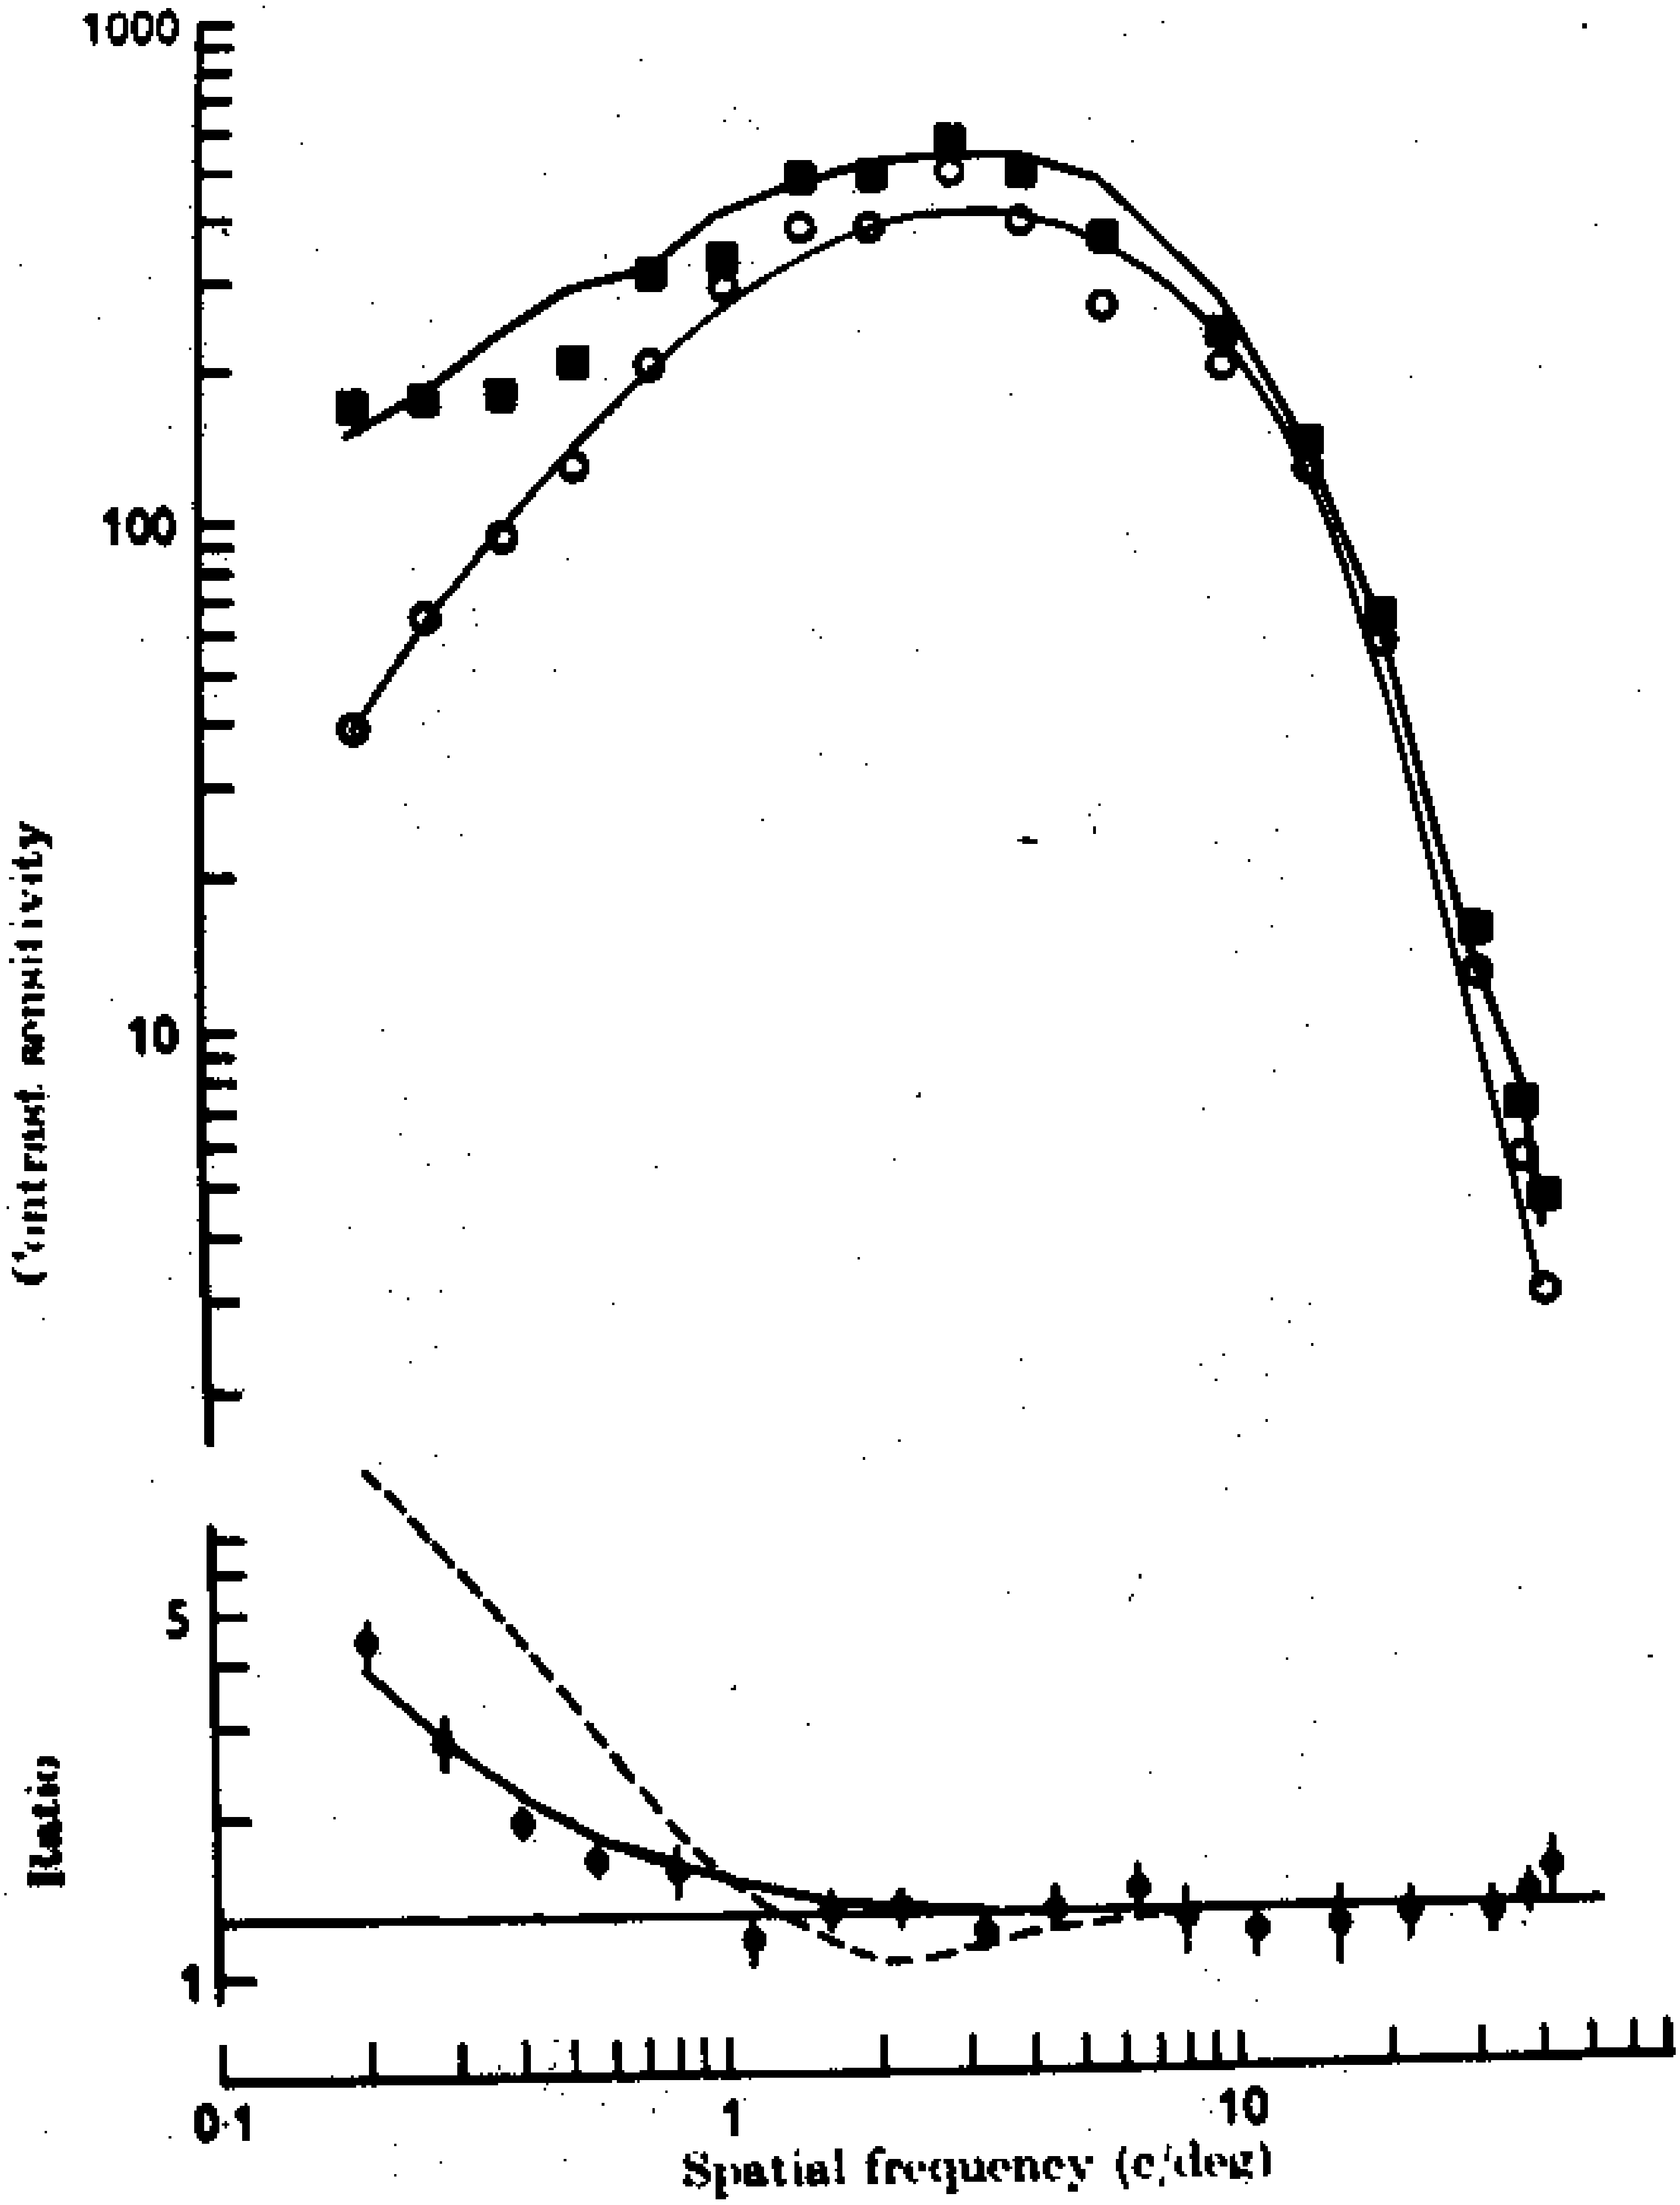
1000
100
10
Contrast sensitivity
5
1
Ratio
0·1
1
10
Spatial frequency (c/deg)

FIGURE 7

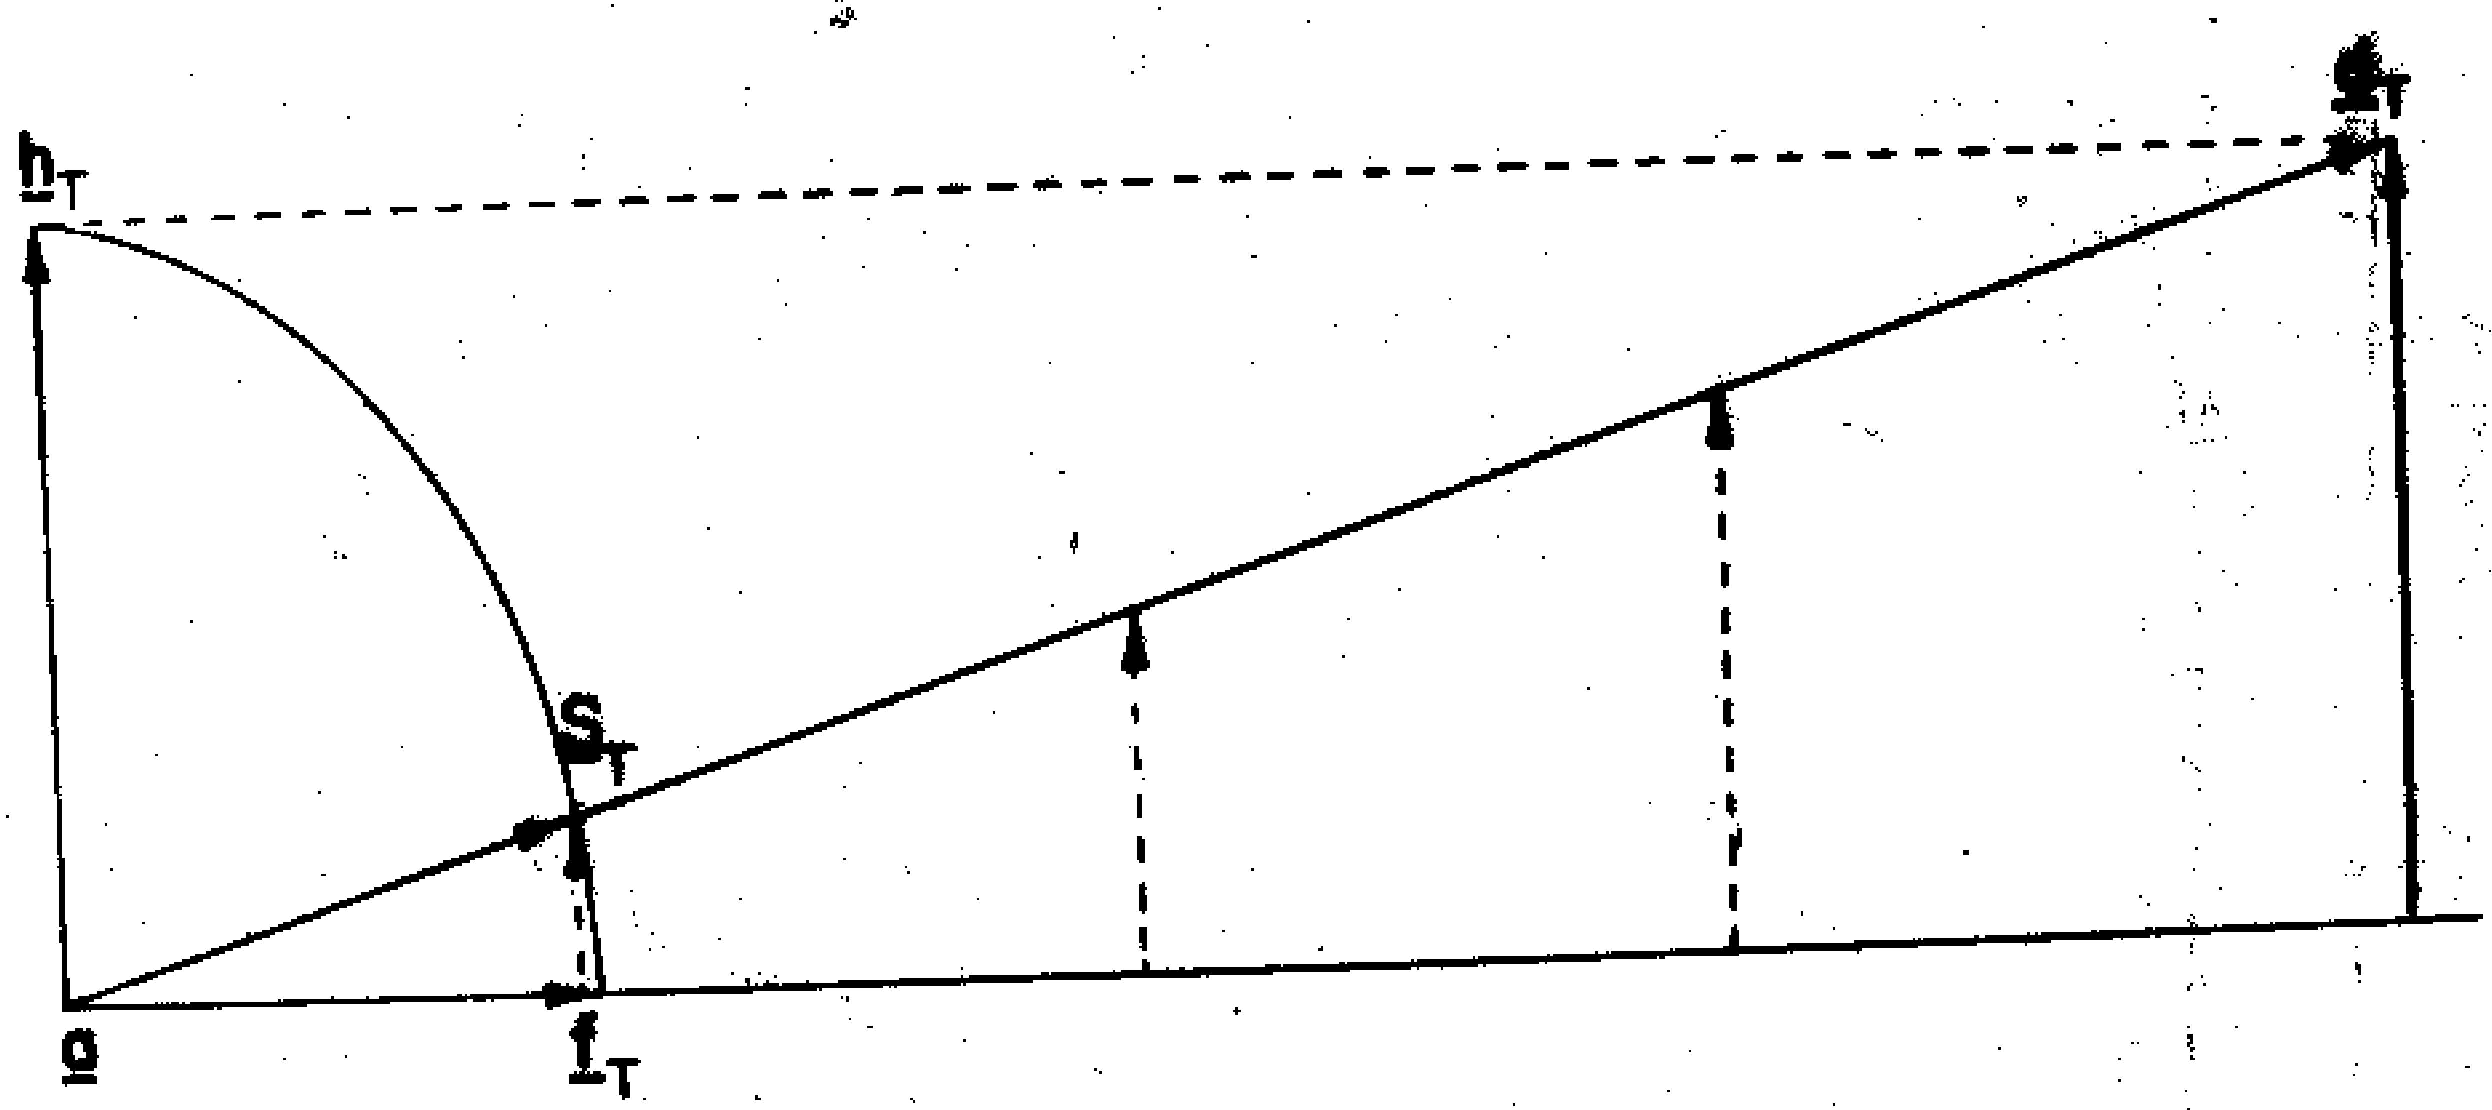

REFERENCES

1. G. Holton, Thematic origins of scientific thought (Harvard, 1973).
2. P. E. King-Smith & J. J. Kulikowski, J. Physiol.,Lond. 247, 237, (1975).
   D. J. Tolhurst, Vision Res, 12, 797, (1972)
3. F. W. Campbell, & J. G. Robson, J. Physiol., Lond. 197, 551, (1968).
4. N. Graham, & J. Nachmias, Vision Res.,11, 251, (1971).
5. A. Pantle & R. Sekuler, Science, 162, 1146 (1968).
   D. A. Pollen, J. R. Lee, & J. H. Taylor, Science, 173, 74, (1971)
   see R. Sekuler, Ann.Rev. Psych., 25, 195,(1974), for a review and bibliography.
6. Similar "feature detectors" have been proposed for patterns with temporal components, e.g., B. G. Breitmeyer & L. Ganz, Psych. Rev., 83, 1, (1976).
   see R. Sekuler in Handbook of Perception, E. Carterette & M. Friedman, (Eds.). (Academic Press, 1974) for a review.
7. See D. Marr & T. Poggio, Science, 194, 285, (1976) for a similar view with respect to stereo disparity detection.
8. D. Middleton, IRE-PGIT, IT-3, 86, (1957).
   D. Middleton. Introduction to statistical communication theory. (McGraw-Hill, 1960).
   C. W. Helstrom. Statistical theory of signal detection. (Pergamon, 1960)
   L. H. Koopmans. The spectral analysis of time series. (Academic, 1974).
9. D. M. Green & J. A. Swets. Signal detection theory in psychophysics. (Wiley, 1966).
   R. Crowzy. In Visual science. J. R. Pierce & J. R. Levene, (Eds.). (Indiana, 1971).
10. L. Padulo, & M. A. Arbib. System theory. (Saunders, 1974).
    A. B. Baggeroer. State variables and communication theory. (MIT, 1970).

10. (Cont'd.)

E. C. Zeeman. In Towards a theoretical biology, Vol. 4. C. H. Waddington, (Ed.). (Eddinburgh, 1972). The sources in the last three notes and in 15 and 16 offer introductions to most of the notions used in this article. The notion that color vision could be described as a manifold dates back to E. Schrödinger. Ann.Physik., 63, (1920). The notion of brainstates as points (attractors) on a high dimensional manifold is probably original with E. C. Zeeman, in Topology of 3-manifolds. M. K. Fort, (Ed.). (Prentice-Hall, 1962). Somewhat more accessible is E. G. Zeeman in Mathematics and computer science in biology and medicine. (M.R.C., 1965).

11. J. Y. Lettvin, H. R. Maturana, W. S. McCulloch, & W. H. Pitts. Proc. IRE, 47, 1940, (1959).

12. e. g., D. H. Hubel, & T. N. Wiesel. J. Physiol., Lond., 160, 160, (1962); ibid., 195, 215, (1968).

13. see H. R. Blackwell, JOSA, 53, 129, (1963).

14. The simplifying assumptions made at a number of points in this paper are not intended to deny that more precise analysis is needed in many cases, but are made to ease explication of the major themes. That is why most of the analysis concerns one-eyed, color-blind perception.

15. S. Lang. Real analysis. (Addison-Wesley, 1969).
M. Spivak. Calculus on manifolds. (W. A. Benjamin, 1965).

16. B. W. Lindgren. Statistical theory, (2nd ed.). (Macmillan, 1968).
D. F. Morrison. Multivariate statistical methods, (2nd ed.). (McGraw-Hill, 1976).

17. Fig. 2-a is from T. E. Cohn, & D. J. Lasley. Science, 192, 561, (1976).

17.(Cont'd.)

Figure 2-b is from D. L. McAdams. JOSA, 32, 247, (1942), reproduced in Y. LeGrand. Light, color and vision. (Chapman & Hall, 1968).

Y. LeGrand. In Visual science. J. R. Pierce, & J. R. Levene, (Eds.). (Indiana, 1971), hints (p. 316) at a statistical explanation for ellipsoidal threshold contours but goes on to reject them.

Figure 2-c is from C. Rashbass. J. Physiol., Lond., 210, 165, (1970).

18. G. Werner, & V. B. Mountcastle. J. Neurophys., 28, 359, (1965), provides evidence that noise is independent of firing rate in cutaneous afferent fibers.

19. Figure 4 is modified by G. Westheimer for Handbook of sensory physiology, Vol. VII/4 Visual psychophysics. D. Jameson, & L. M. Hurvich, (Eds.). (Springer-Verlag, 1972), from F. L. Van Nes & M. A. Bowman, JOSA, 57, 401, (1976). This latter paper offers much grist for further analysis of the relationships.

20. G. Hadley. Linear algebra. (Addison-Wesley, 1961).

21. C. Blakemore & F. W. Campbell. J. Physiol., Lond., 210, 165, (1969), and A. Pantle, & R. Sekuler, ibid.

22. R. E. Armstrong. Doctoral dissertation, in progress.

23. J. Nachmias, & E. Kocker. JOSA, 60, 382, (1970).

W. P. Tanner. Ann. N.Y. Acad. Sci., 89, 752, (1961).

24. M. Sachs, J. Nachmias, J. G. Robson. JOSA, 61, 1176, 1971).

25. B. Scharf. In Foundations of modern auditory theory. (Academic, 1970).

26. C. F. Stromeyer, III, & B. Julesz. JOSA, 62, 1221, (1972).

27. The rather long incubation period of the ideas in this paper were supported in part by NIH grant Ey-321 to Robert Sekuler. More recent support has been mainly supplied by my wife, Cyn, and in the final stages of preparation by Donald T. Campbell. Of all those who contributed to the creation of the paper, I wish to thank especially Robert Williams who taught me the legitimacy and utility of abstract and bizarre thought.