PATTERN RECOGNITION AND THE RESPONSES OF SENSORY NEURONS*

H. B. Barlow

Neurosensory Laboratory, School of Optometry
University of California, Berkeley, Calif.

In 1959, Lettvin and his associates announced to the electrical engineering world that the five classes of retinal ganglion cell in the frog each responded to a different, specific, spatiotemporal pattern of excitation. These cells were, in fact, performing "feature extraction," which at that time was coming to be recognized as a necessary first step in more complex pattern recognition. I remember Keffer Hartline telling me of these claims at about the time they appeared. Both of us had worked on retinal ganglion cells in the frog, and at first I don't think either of us believed much of what Lettvin and his co-workers reported. Hartline had worked on movement sensitivity and was particularly surprised by "unidirectional rate sensitivity"—the claim that some ganglion cells responded to movement in one direction and not the other.

At first I was also surprised by this, but uneasy memories came back of some observations I had made on *on-off* units a few years earlier. I had been reading Sperry's experiments on optic nerve regeneration in frogs whose eyes had been excised and rotated, and I was very impressed by the incidentally mentioned fact that, in the normal frog, the left eye only evokes optokinetic movements to the right, and vice versa. I was puzzled by the movement sensitivity of *on-off* units at the time, and I looked to see if this optokinetic behavior was of retinal origin: Did the ganglion cells of one eye only respond to movement in the appropriate direction? I looked long enough and hard enough to convince myself that my hypothesis was wrong, for the ganglion cells certainly responded to image movements that would have failed to evoke optokinetic movements in the intact animal. But when Lettvin jogged my memory, I recalled that I had sometimes obtained "funny" results: Some ganglion cells did seem to show a preference for movements in a certain direction, but this direction had not usually corresponded to the direction required by my hypothesis. Beyond using these facts to throw out the hypothesis, I did not pay much attention to them, and I did not pursue the matter because it was so difficult to produce smooth movements in a variety of directions with the apparatus that was available. In retrospect, I think I only once or twice saw a degree of asymmetry of response that I would now accept as proof of directional selectivity, but I missed it then and only became fully convinced that retinal ganglion cells can be directionally selective when Hill and I observed this phenomenon in the rabbit (Barlow & Hill, 1963). Since then, I think we have found out something about the mechanism by which the nervous system abstracts patterns, and I shall return to

---

*This work was supported in part by United States Public Health Services research grant no. NB-05215.

this, but first I want to explain another reason why I was unreceptive to directional selectivity when I first had the opportunity of investigating it.

Pattern Recognition and Noise

It seems to me that the central puzzle about pattern recognition is not the ability to respond to one specific pattern, but the ability *not* to respond accidentally or spuriously to incidental stimuli that contain some features of the specific pattern. To put this point another way, there is no difficulty about pattern recognition in the absence of noise, but there are great difficulties in the presence of noise. For this reason, I was fixated on the difficulty and importance of signal/noise discrimination in sensory systems. This is a problem that cannot be separated from pattern recognition, inasmuch as you cannot get good discrimination without knowing what pattern you are looking for; on the other hand, you cannot deal with pattern recognition prior to the signal/noise discrimination, for then you would only obtain good performance on the necessarily limited set of patterns you initially looked for. If one is interested in signal/noise discrimination, one naturally looks at the quantitative aspects of excitation by light, and I was therefore more interested in things like Weber's law and Ricco's law of areal summation at the time when I had the opportunity to look at movement sensitivity in the frog.

The detection of small quantitative differences in light excitation and the use of this information in identifying trigger features, such as movement in a particular direction, are fascinating aspects of the integration of sensory information. Perhaps the biological tricks used in this task would be of interest to pattern-recognition programmers, for it seems to be in the combination of many weak cues against a background of noise that human pattern recognition outperforms machines.

Invariance of Response to Trigger Features

Another aspect of the same problem is the capacity to maintain the same differential response to movement in spite of very great variations in stimulus conditions, as shown in FIGURES 1 and 2. These were prepared by W. R. Levick, and they prove another of Lettvin's points, namely, the wide range of invariance for detection of a particular pattern. The response does not depend upon the light intensity, or even the direction of the contrast; on the other hand, the information about intensity is required for performing the discrimination. I cannot conceive of any mechanism capable of this performance that does not require adjustments when the background luminance varies. What we learn from this kind of result is not only that these adjustments must be made, but also that they must be made at a very early level, probably in the receptors themselves if the retinal pattern recognition mechanism I shall describe is correct.

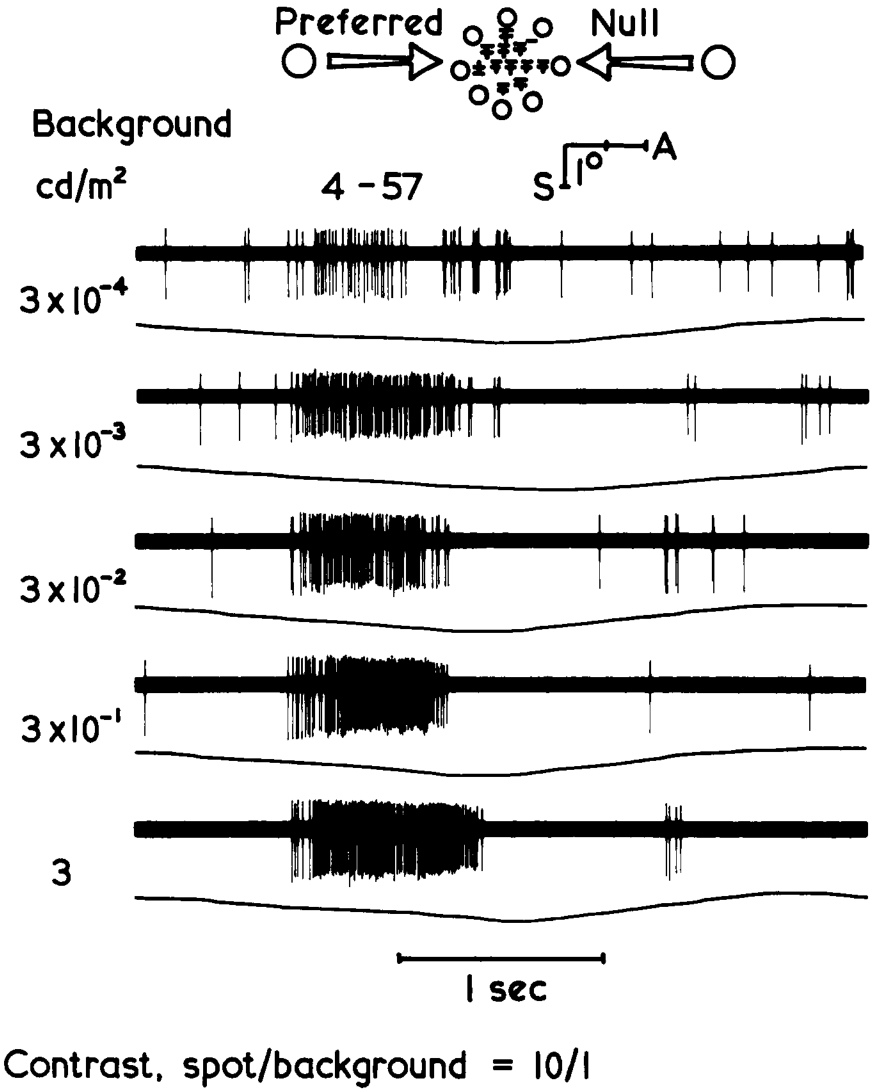

FIGURE 1. Invariance for changes of background intensity of the directionally selective responses of a rabbit retinal ganglion cell. A spot with ten times the background luminance was moved across the field in the preferred direction and then back in the null direction. The line below the action potentials indicates the position of the spot. The response is more vigorous at high backgrounds, but the directional preference is the same. This illustration was prepared by W. R. Levick.

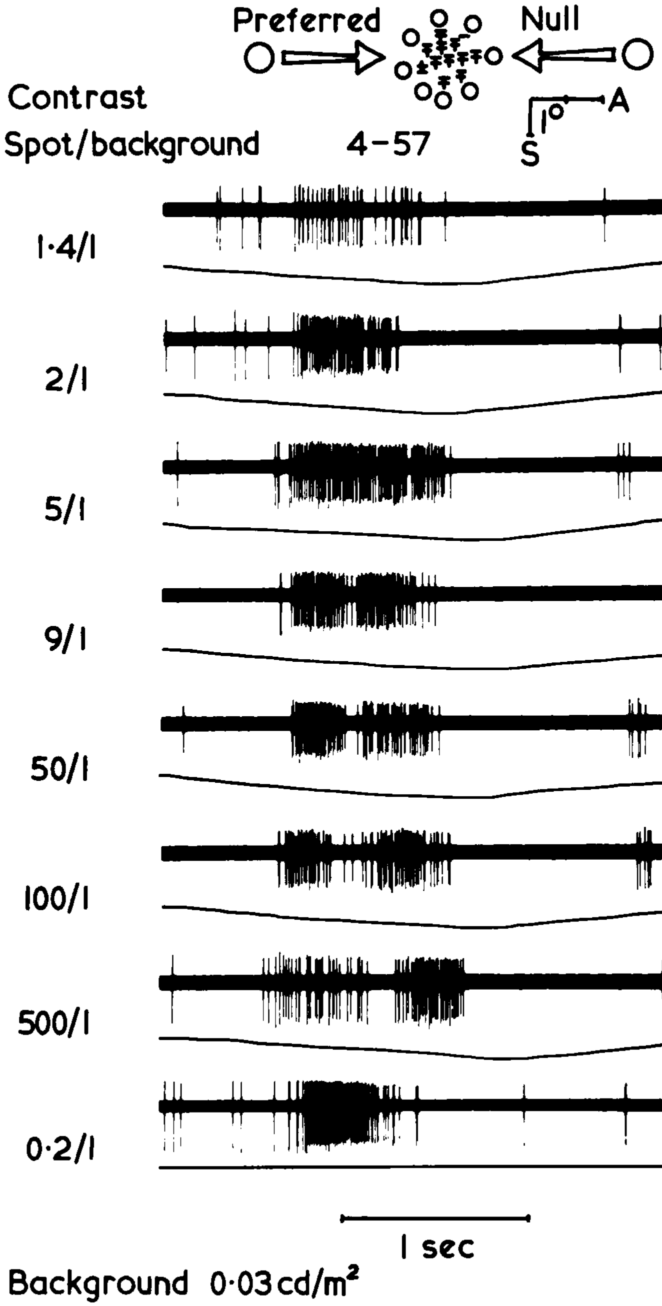

FIGURE 2. Invariance for changes of spot luminance and contrast of the directionally selective responses of a rabbit retinal ganglion cell. As in FIGURE 1, the same directional preference is obtained at all luminances, and even with contrast reversed. This illustration was prepared by W. R. Levick.

Neural Mechanisms of Pattern Recognition

FIGURE 3 shows diagrammatically the various mechanisms that have been proposed for neural pattern recognition. Lettvin and his associates (1961) believe that the anatomical form of the retinal ganglion cells, tracings of which by Ramon y Cajal are shown here, reveals or reflects the pattern of connections they make to the bipolar cell endings in the inner plexiform layer. They think that different levels in this layer represent particular functional derivatives of the pattern of light falling on the receptors, and that a ganglion cell can thus form the appropriate combination of these functions by dipping down to the right depth at various points.

The anatomy is quite different in the scheme of Hubel and Wiesel (1962, 1965); they have evidence that the cells in the cortex which connect to a higher-order cell are grouped in columns perpendicular to the cortical surface. However, the principle whereby specificity is achieved is the same as in Lettvin's scheme, and might be called "selective summation," i.e., by connecting to a specific subset of cells at a previous level, a high-level cell comes to respond only to a particular configuration of excitation at the lower level, e.g. that which results from presentation of a black bar at a particular angle in one position in the visual field. Since these cells respond very poorly to part of the black bar but well to the whole of it, selective summation in this case corresponds roughly to logical conjunction—both parts of the bar are required for a vigorous response. However, Hubel and Wiesel suggest that selective summation also performs the equivalent of a logical union: If a high-level cell connects to many cells responsive to bars at one and the same orientation, but at different positions in the visual field, then this cell will generalize for position and respond to a bar at the correct orientation regardless of position. Thus, we have a possible mechanism for position invariance.

Levick and I were led to the third scheme to account for the directional selectivity of the rabbit retina. We postulate that specificity is achieved by connections from an inhibitory associational neuron that prevents the cell from responding to stimuli that do not possess all the specific requirements of the trigger feature. This inhibitory or vetoing cell would play a key role in preventing unwanted responses of the pattern recognizer when it is confronted with a noisy input containing some, but not all, of the features of the desired pattern. In addition, this scheme, unlike that of Hubel and Wiesel, provides different mechanisms for generalizing a stimulus and for achieving the required specificity—two operations that are logically quite different.

One can express the operation of this mechanism in logical terms as follows: If E^{ψ} is the class of inputs which fires a neuron so situated that ψ synapses are passed before input activity reaches it, and if $E^{\psi+1}$ is the class for a neuron after one more synapse, then, using v to signify "or" and . ~ for "and not,"

$$E^{\psi+1} = (E^{\psi}_{a} \vee E^{\psi}_{b} \vee E^{\psi}_{c} \ldots) . \sim (E^{\psi}_{r} \vee E^{\psi}_{s} \vee E^{\psi}_{t} \ldots) \qquad (1)$$

The sensory stimuli we tend to class together seem to have great generality com-

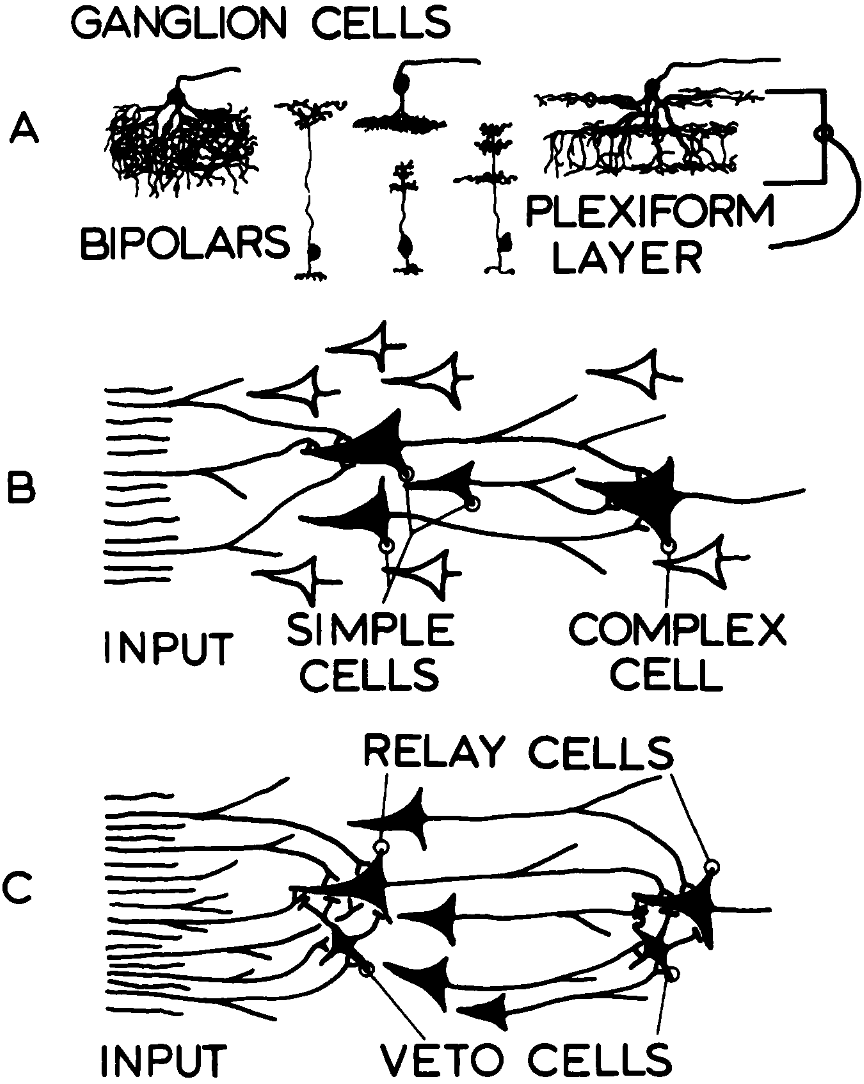

FIGURE 3. Three suggested schemes for neural pattern recognition. (A) Lettvin and coworkers (1961) suggest that frog retinal ganglion cells dip their dendrites to various depths in the internal plexiform layer, and form a combination of the functional derivatives of the input which are relayed to these layers by the bipolar cells. (B) Hubel and Wiesel (1962) suggest that "simple" cortical neurons in the cat sum from a selected subset of the input from the lateral geniculate, "complex" cells sum selectively from the "simple" cells, and so on. (C) Barlow and Levick (1965) suggest that generalization, or position invariance, is achieved by selective summation, as in part B, but that specificity of response is aided by associational cells that veto or inhibit unwanted responses. Repeated application of these operations in a hierarchical system would be required to recognize and discriminate complex patterns.

bined with a tantalizing specificity, and Equation (1) has the merit of expressing this. It also implies that the mechanism is hierarchical, which is realistic. However, we do not know if it is correct in other ways; certainly it is inadequate in one respect: Like all logical schemes, it implies that the variables are binary, whereas we know that they are not. We come back again to the problem of how the intensive aspects of light stimuli are handled, or, as shown in FIGURES 1 and 2, how they are deliberately ignored and discounted.

Light Intensities and Certainty Signals

FIGURES 4 and 5 offer a suggestion about this problem. The absolute intensity of light is often not important in recognizing the kind of patterns to which biological systems respond: Certain sensory events have occurred, and what one optimally needs are the probabilities of these events on certain hyptheses. For instance, if I enter my darkened home and come to the point where there is a swing door, I must decide whether to put out my hand and push it open or whether to walk straight on. The decision should be made partly on the basis of what I can see, but this information must be combined with the prior probabilities of the door being open or closed, and also the payoffs and penalties of the various outcomes; the survival value of walking through a closed door is obviously lower than it would be for the other possibilities. However, these are not sensory problems; what we need from our eyes is simply the probability of the sensory events occurring if the door is closed, and so on.

Let us suppose that the door is white and that the area we would see through the door if it were open is black. To decide whether to raise my hand and push, I need to know the probability of the sensory events that occurred in the retinal region for the door *if* the light intensity in that region is zero, and the corresponding information for other hypotheses. Actual figures for the numbers of quanta absorbed, or estimates of light intensity, if they were available, could form the basis for computing these probabilities, but it is only the probabilities themselves that can be combined directly with the information on payoffs and prior probabilities to form an optimum basis for rational decision.

This line of reasoning applies also to the utilization of weak and remote cues in pattern recognition. To decide what letter separates *c* and *t* in a three-letter word, we can establish probabilities from our knowledge of possible words and from the more remote context of the sentence. What we require from the region of the letter itself are probabilities: What is the probability of this particular pattern of blackening if an *a* was intended, or an *o*, or a *u*, and so on? I think it is now clear that the nervous system does far more sophisticated pattern recognition in the periphery than would have been suspected ten years ago. Is it also more sophisticated in the quantitative expression of the result? Is a single unit signaling, not just the presence of its trigger feature but the degree of certainty that the trigger feature is present?

From this point of view, *on* and *off* receptors are signaling how certain it is that a change of light has occurred in the direction of their predetermined sensitivity (i.e., to an increase or decrease of light). The magnitude of the signal is related to the magnitude of the change, but it would be an oversimplification to regard it as a function of light intensity alone. Light adaptation and lateral inhibition can readily be interpreted as normalizing operations on this basis, but I cannot argue that point in this paper. Instead, I want to offer some evidence in support of the notion.

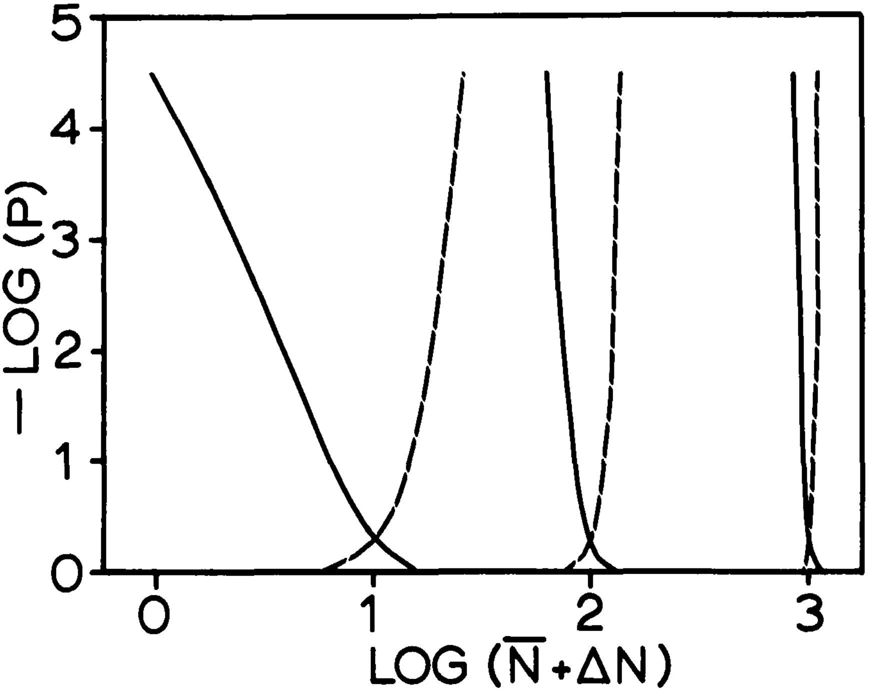

FIGURE 4. Certainty signals from idealized *on* and *off* units. The three sets of intersecting lines refer to the hypothesized conditions that the unit is adapted to the absorption of an average of, respectively, 10, 100, or 1000 quanta per unit of time. The light is then assumed to be suddenly changed to values leading to a mean rate of absorption shown on the horizontal axis. The probability (p) of the new rate of quantal absorption occurring *if* the mean rate did not change is calculated, and $-\log p$ is plotted vertically. This is a measure of the certainty that the light was, in fact, turned *on*, or *off*, derived from the sensory events available to the certainty signaling unit.

FIGURE 4 shows calculated relations between impulse frequency and intensity for idealized *on* units and *off* units. The trigger feature for *on* is an increase in light intensity. To obtain a measure of the certainty that this has occurred, one goes through the same maneuvers as a statistician. First, the negation of the trigger feature becomes the null hypothesis: "There has been no increase in intensity." Then, the probability according to the null hypothesis of the events that have occurred is calculated. Finally, a low value of this probability is taken as a high degree of certainty that the trigger feature was present. In FIGURE 4, the negative logarithm of the probability on the null hypothesis is plotted as a measure of certainty; this choice of function is arbitrary, but it has certain obvious advantages when several probabilities are to be combined. The horizontal scale represents log (average number of quanta absorbed); it is assumed that this number is initially either 10, 100, or 1000, and then changes to a neighboring value. From cumulative Poisson tables, or gaussian approximations, the probability of the new value occurring at the old mean rate is obtained, and its negative log is plotted vertically. Broken lines give these values for *on* units; solid lines show the values for *off* units.

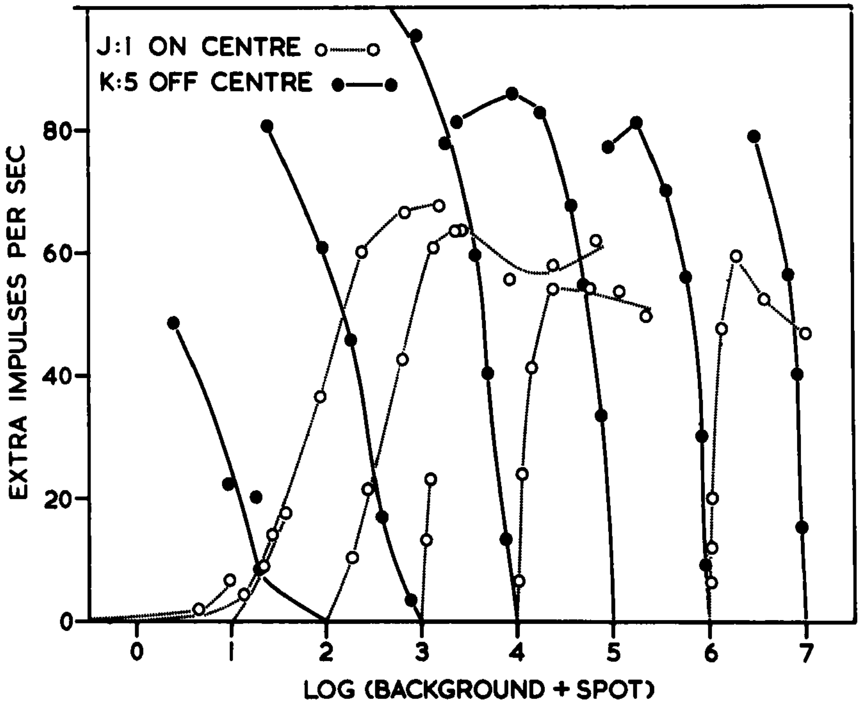

FIGURE 5. Cat retinal ganglion cells as certainty signalers. The responses of an *on*- and an *off*-center unit were measured for a light turned *on* and *off* for 1.28 seconds. The luminance was held at an adaptation level indicated by the points where the lines cut the horizontal axis. A spot on the center was turned *on* and *off*, and responses at *on* or *off* are plotted in impulses per second above the luminance resulting from spot and background together. These results have some features in common with those calculated in FIGURE 4 for an ideal "certainty signaler."

FIGURE 5 shows how Levick and I found that the *on*- and *off*-center units of the cat behaved after they were adapted to various steady levels and then the intensity in the central area was stepped to a new, adjacent value for 1.28 seconds. For the *on* units, an increased number of impulses was evoked, and the average increase in rate is plotted as impulses per second vertically, with total luminance of spot and background plotted horizontally. The point where these lines cut the zero impulses per second axis is the adaptation level. For *off* units, the excess impulses occurred when the light was turned off, and they were counted for the 1.28-second *off* period. Unfortunately, this illustration was made to show the incremental number of impulses above the maintained discharge, and it therefore does not show the decrement at *off* in the *on* unit, and vice versa.

There are some similarities and some differences between the hypothetical "degree of certainty" messages of FIGURE 4 and those actually evoked in FIGURE 5. Note that *on* and *off* messages acquire a meaningful complementarity on this scheme, and that the lines become steeper at higher luminances. One of the main differences is that the real messages saturate; this occurs at what is physiologically rather a low rate of about 100 impulses per second, but this was the mean rate over 1.28 seconds, and the peak rate was, of course, much higher.

Conclusion

Sensory neurones are quite sophisticated pattern recognizers, and their graded responses may signal the certainty with which the events in the receptors indicate that the trigger feature is present.

References

BARLOW, H. B. & R. M. HILL. 1963. Selective sensitivity to direction of motion in ganglion cells of the rabbit's retina. Science **139:** 412–414.

BARLOW, H. B. & W. R. LEVICK. 1965. The mechanism of directionally selective units in rabbit's retina. J. Physiol. (London) **178:** 477–504.

HUBEL, D. H. & T. N. WIESEL. 1962. Receptive fields, binocular interaction, and functional architecture in the cat's visual cortex. J. Physiol. (London) **160:** 106–154.

HUBEL, D. H. & T. N. WIESEL. 1965. Receptive fields and functional architecture in two non-striate visual areas (18 and 19) of the cat. J. Neurophysiol. **28:** 229–289.

LETTVIN, J. Y., H. R. MATURANA, W. S. McCULLOCH & W. H. PITTS. 1959. What the frog's eye tells the frog's brain. Proc. Inst. Radio Engr. **47:** 1940–1951.

LETTVIN, J. Y., H. R. MATURANA, W. H. PITTS & W. S. McCULLOCH. 1961. Two remarks on the visual system of the frog. *In* Sensory Communication. W. Rosenblith, Ed.: 757–776. (Massachusetts Institute of Technology Press) John Wiley & Sons. New York, N.Y.