

Coding in the Auditory Nervous System

by
I. C. WHITFIELD

Neurocommunications Research Unit,
University of Birmingham

Experiments on how hearing works suggest that it is the connectivity of the central pathways, and not the existence of elaborate code patterns in individual structures, which mediates discrimination of such a wide range of sounds by the ear.

THE idea of auditory analysis by resonant structures is one of some antiquity. The earliest theories were framed in terms of cavity resonances, for the inner ear was then thought to be filled with air. Quite explicit theories of this type are to be found at the beginning of the seventeenth century. From the time of du Verney¹ onwards, however, the solid structures to which the auditory nerve is distributed came to be regarded as the resonators, and the only changes between 1700 and 1850 resulted from a continuing search for the precise location of the resonators, the field of search being progressively opened up by improvements in microscopic and histological technique. The postulated resonance seems to have been regarded mainly as a means of magnifying the weak sound stimulus rather than as an analytic mechanism, although du Verney did consider that as a result of differential vibration the auditory nerve would "receive different impressions which represent in the brain the various characteristics of tones". Haller² appears to have been the first to speculate on the possibility of distinguishing sound frequencies by the frequency of the 'tremors' in the auditory nerve fibres and thus to have foreshadowed the other great group of auditory coding theories—the frequency or "telephone" theories first formalized by Rutherford in 1886.

In 1857 Helmholtz³ produced his famous theory, which can probably be regarded as the first attempt to formulate a detailed system of auditory coding. He did this by combining three previously enunciated principles. The first of these was Ohm's law of auditory analysis. It is a matter of common experience that, in notes sounded by most musical instruments, it is possible with practice to distinguish both the fundamental and one or two overtones. Ohm extended this observation in his hypothesis, and stated that any complex periodic sound wave behaved—as far as the ear was concerned—as if it consisted of a series of suitably related sine waves corresponding to its Fourier components.

Helmholtz was quick to recognize that a system of tuned resonators is precisely the kind of structure to realize Ohm's proposal physically, and he therefore emphasized the analytic rather than the magnifying properties of the resonance hypothesis of cochlear function already established.

The third principle was that which is now commonly known under the name of Müller's doctrine of specific nerve energies. Müller proposed, in effect, that there are five different kinds of sensory nerves—one for each of the five senses—and that stimulation of one particular kind of nerve, by whatever means, evokes specifically the sensory quality to which it is related. Other people carried the subdivision further and the idea reached its fullest extension with Helmholtz's proposal of a specific fibre or group of fibres for each distinguishable pitch. The Helmholtz theory, in its original form, was a pure place theory. Each subjectively distinguishable tone

had a corresponding resonator in the cochlea and each resonator was supplied by an individual nerve fibre or group of fibres distinct from all others. Frequency was thus coded in terms of activity or inactivity in particular fibres, and intensity in terms of the degree of that activity.

There are a great many problems raised by the Helmholtz hypothesis, of most of which Helmholtz was himself aware. It is not possible to discuss them in detail here; they have been well treated by Wever⁴. It is sufficient to point out the two general headings: the identity of the resonators, and their selectivity. In spite of the greatly increased detail which the compound microscope had revealed by the mid-nineteenth century—in particular the structure of the organ of Corti—it never proved possible to identify any component which would account in number and range of physical properties for the number and range of the known discriminable frequencies. Likewise, it proved impossible to reconcile the high degree of selectivity called for in the resonators with the experimentally determined decay time of the system.

These two difficulties caused Helmholtz later to abandon the idea of highly selective resonators and to suggest only that one particular resonator was most highly excited by a tone while those for some distance on each side would be excited to lesser extents. He thus implicitly abandoned the specific nerve energy "one tone, one fibre" code in favour of a code in which there was a non-unique array of fibres subjected to a unique pattern of activation.

Although both the general principle and the details of resonance in the cochlea continued to be matters of argument, the theory was quite widely accepted and taught in elementary textbooks for some seventy years. Perhaps its very elegance and simplicity made people feel it ought to be true and encouraged its acceptance.

Behaviour of the Cochlear System

The essence of the Helmholtz theory was a series of transverse mechanical resonators, distributed along the length of a tapering basilar membrane. Tones of different pitch would set into vibration different but highly localized regions along the length of this membrane. Until the early 1940s no one had actually looked at the basilar membrane during activation of the ear by a sound stimulus. Then Békésy⁵, who had been working for some time on cochlear models, published several papers on the mechanics of real ears. The essence of his method was to examine the vibrating membrane through an artificial window in the cochlea with the aid of stroboscopic illumination. He found indeed that there was a region of maximum vibration of the membrane for a given frequency of stimulus, and that this maximum moved from the apex towards the base of the cochlea as the sound frequency was increased. The maxima were not very sharp, however, and considerable overlap of the vibration envelopes was evident between frequencies an octave or more apart. The method also made it possible to determine the

ive phase of the vibration, and the result proved against a resonance theory. The phase lag of placement of the basilar membrane relative to the driving force increases progressively from the base to the apex of the cochlea, reaching values as large as 5π at the apical end. The lag at the position of maximum displacement is 2π or more. These values are not compatible with a simple resonant system. The disturbance can best be thought of as a travelling wave originating at the basal end of the membrane and travelling towards the apex of the cochlea. As it travels it gradually increases in amplitude and then dies away rather abruptly. With high frequency driving forces the peak amplitude is reached close to the basal end, and the wave does not travel very far before dying out. With lower driving frequencies the wave travels farther and farther along the cochlea, and the maximum is progressively displaced towards the apical end (Fig. 1).




Fig. 1. "Instantaneous" pictures of the travelling wave on the basilar membrane for a high frequency, a; a middle frequency, b; and a low frequency stimulus, c. The dashed envelopes show how the disturbance travels farther along the membrane from base to apex before reaching a peak and dying out.

The basilar membrane carries on its surface a rather complex structure known collectively as the organ of Corti. This structure contains the hair bearing cells which form the mechano-neural transducers. Vibration of the basilar membrane results in deflexion of the hairs attached to these cells and this, in turn, by a process not fully understood, results in a discharge of impulses in the sensory nerve fibres associated with these cells. The greater the amplitude of vibration of the membrane, the faster are impulses discharged by the nerve.

Although the innervation of the cochlea is quite complex, in general terms we can say that the 30,000 to 40,000 fibres of the auditory nerve are evenly distributed to hair cells along the whole length of the basilar membrane. Because the vibration envelope of one tone can overlap that of other tones an octave or more away, and because the extent of the envelope varies with intensity, we might expect that a single auditory nerve fibre could be activated by many different tones, and this proves to be the case. Near its threshold a fibre responds to only a small range of frequency, but as the sound intensity is increased it responds over a greater and greater range until at high intensities it may be responding to any frequency within a band several octaves wide. Each

fibre thus has a triangular frequency/intensity response area which roughly mirrors the shape of the vibration envelope.

Because any one fibre responds to a wide band of frequencies, it would be expected, conversely, that a large number of auditory nerve fibres will be discharging at any given stimulus frequency. Schuknecht⁷ has shown that at the quite moderate sound pressure level of 40 dB, between 15 and 25 per cent of the whole system is activated by any one single frequency. But 25 per cent of 40,000 fibres is 10,000 fibres. It follows that if the frequency is changed by 1 per cent—an easily detectable change—something like 200 fibres will be added to one end and a similar number subtracted from the other. In other words, 98 per cent of the active fibres are activated by both tones. When we remember that the total number of fibres activated depends on intensity as well as frequency along the array, it is clear that there is no question of there being specific groups of fibres for specific discriminable frequencies. How then does the system function?

Possible Pulse Interval Codes

An alternative to the resonance hypothesis was put forward by Rutherford⁸ in the 1890s and has become known as the "telephone theory". Rutherford suggested that, rather than peripheral analysis taking place, the auditory nerve transmitted a frequency replica of the sound. Little was known in those days about the properties of the nerve impulse, but subsequent work has shown that it is not possible for a nerve to carry trains of impulses at frequencies greater than at most 1,000 per sec, and it usually cannot sustain rates greater than 500/sec for more than a very short time. This leaves a large proportion of the audible spectrum without representation. To deal with the frequencies above 500 c/s, it has been suggested that some form of alternation takes place. According to this idea, a frequency of 10,000 c/s would be dealt with by twenty fibres each firing 500 pulses/sec and firing in rotation. Of course, this would only signal the frequency element of the sound, and considerably more fibres would be required to deal with the intensity parameter, but there are some 40,000 fibres available so that the hypothesis does not immediately fail on this score.

The primary question, however, is the experimental one: do the pulses in auditory nerve fibres bear any relation in time to the stimulus frequency? At low stimulus frequencies (say, 200 c/s) there is undoubtedly a very strong relationship. However, there is a good deal of "jitter" in the system which causes this relationship to become progressively less and less obvious as the stimulus frequency is increased. In the auditory nerve this relationship fails somewhere between 2,000 and 4,000 c/s, according to the time and space over which one chooses to integrate the measurements. The situation gets progressively worse as one ascends the neural pathway, owing to the variable nature of synaptic delays (the variation may be as much as 300 μ sec). This means, in effect, that for anything but the very lowest frequencies, a frequency code would have to be converted to some other form at a very early stage in the neural path. There is no evidence that any such conversion takes place. We do not, for example, find that single neurones at the higher levels of the nervous pathway respond only to narrow frequency bands; the band of frequencies to which a given fibre responds at, say, the collicular level is just as great as that of an auditory nerve fibre. We must conclude that although frequency representation is a possible mechanism for low frequencies, it is not a valid one for high frequencies.

Even at low frequencies there is a further problem to be considered. It is very easy when recording nerve impulses on an oscilloscope, or using them to prepare a pulse interval histogram, to distinguish a frequency of 990 from a frequency of 1,000. However, for the nervous system to do this, it too must have a suitable clock and there is no

evidence that a neural clock of this degree of accuracy exists.

We have considered two types of coding—individual channel activity and mean pulse frequency. Let us now consider another. The pulse trains in a nerve fibre are not, in detail, evenly spaced in time. This is especially true at higher levels in the system. It is clearly possible that a pulse interval code could exist, and such a code has a potentially high information capacity. In the past ten years or so, the possibility of such a code in the nervous system has attracted a lot of attention and a great deal of work has been done on analyses. So far, little encouragement has been forthcoming. Examination of such "patterns" in the auditory system has served only to support the conclusion that the distributions of pulse intervals are random.

In the early stages of the system; if we discount the relative deficiency of short intervals imposed by the refractory properties of the excitable elements, and any residual first order patterning of the type referred to here, there remains only an interval distribution which could adequately be accounted for by a random process. At higher neural levels the distributions are complicated by multiple delay paths, but there still appears nothing obviously specific to the stimulus. The objection raised here would also seem to apply even more cogently to the use of a pulse interval code: is there in the nervous system any mechanism for decoding such a pattern? Again we know of none.

Coding of Stimulus Frequency and Intensity

Before going further, it may be worth while to consider what the auditory system actually has to do in hearing, say, a steady tone or a vowel sound. If we make measurements of the frequency and intensity difference limina (DLs) over the whole range of hearing, we can construct a two-dimensional diagram consisting of a number of cells or compartments which show at a glance the size of the DLs in any particular region⁹. If we count up the total number of these cells we find there are about 340,000. However, although many people have greater or lesser degrees of absolute pitch it would be absurd to suggest that anyone when presented with a tone at random can say "That is 6,130 c/s at 30 dB above threshold", and the auditory system is not called on to transmit data for any



Fig. 2. Diagram to show the distribution of activity in the array of fibres originating from the basilar membrane. The height of each line represents the mean rate of nerve impulses in the corresponding fibre. The upper and middle diagrams represent two different stimulus frequencies, while the lowest diagram represents an increased stimulus intensity. In actuality there would be several hundred times as many fibres activated as are shown here.

such choice. Intensity difference limina are normally determined by comparing adjacent stimuli. Again, although speech has a highly complex waveform of potentially very high information content, little of this is used at any one time. Even highly degraded waveforms such as differentiated clipped speech yield 92 per cent intelligibility scores¹⁰. It seems that a very few frequencies, usually called the formant frequencies, or perhaps even the second formant alone, are sufficient to characterize the speech sounds as far as intelligibility is concerned.

Let us see, therefore, how far we can get with a rather simple code. We have already seen that the pattern of activity in the array of auditory nerve fibres for a tonal stimulus is like that of Fig. 2 with fibres firing at various rates according to the degree of excitation of the part of the membrane from which they arise. In this array we could, for example, assess intensity in terms of the discharge rate of the most highly excited fibre, or by the total number of fibres carrying any activity. Likewise, we could assess the stimulus frequency by the position of the most active fibre (the position of the "peak") or by the relative position of the two ends of the array of active fibres. Neither of these two latter criteria is very good, however, as a small percentage of 'noise' pulses added at successive synaptic junctions could easily bias their locations.




Fig. 3. Transformation of pulse rate distribution in the fibre array effected by the cochlear nucleus. *a*, input pattern; *b*, output pattern. Thick lines represent a low stimulus intensity, thinner lines a high intensity. Other conventions as in Fig. 2 (after Whitfield¹¹).

If we examine the behaviour of a single auditory fibre at a higher level beyond the first cell station (cochlear nucleus) we observe two differences from an auditory nerve fibre^{12,13}. First, the discharge rate is no longer a monotonic function of stimulus intensity, so that intensity cannot be signalled purely in terms of discharge rate. Secondly, over a rather wide range of stimulus intensity there is little change in firing rate. The curve tends to increase rather rapidly near threshold as the intensity is raised, and then to flatten off, falling again somewhat at very high intensities. Allanson and Whitfield¹⁴ have pointed out that this effect leads to a transformation in which the noise-sensitive auditory pattern is transformed into one in which all the fibres are substantially either "hard on" or "hard off" (Fig. 3). In other words, if we consider only whether there is activity or no-activity in a given fibre, then the frequency is signalled by the relative position of the two edges, and the intensity by the number of active fibres between those edges. Such an arrangement is not only very resistant to degradation, but presents the information in a form which should be eminently detectable by the nervous system which is well adapted to locating regions of abrupt change in activity.

An immediate difficulty arises. Suppose we have not one single tone as stimulus, but two or more simultaneously as will certainly occur very often in real life. Because of the considerable overlap of activity, the auditory nerve pattern might look something like Fig. 4*a*, and this, after transformation, might be expected to give rise to a single block of active output fibres indistinguishable from that produced by a single, much louder tone, the frequency of which is located somewhere between the two tones we are considering. If we actually try the experiment, we find that although the fibres in the 'overlap' region are activated by either of the two tones sounded alone, if we sound them simultaneously these fibres are actually




Fig. 4. Input/output relations for a stimulus involving two frequencies. Note the inhibition (i) in the "overlap" region. Conventions as in Fig. 2 (after Whitfield¹⁹).

inhibited and fail to respond (Fig. 4b). This type of inhibition is a common feature of sensory pathways^{15,16}. The effect of its existence is clearly to preserve the identity of separate stimuli—in this case the blocks of active fibres corresponding to the two tones.

It is also found experimentally that the fibre array preserves an orderly anatomical arrangement at least as far up as the inferior colliculus¹³. This means that channels which are adjacent frequency-wise are also spatially adjacent. The activity in the total array in response to a complex sound signal will therefore appear rather like that shown diagrammatically in Fig. 5a.

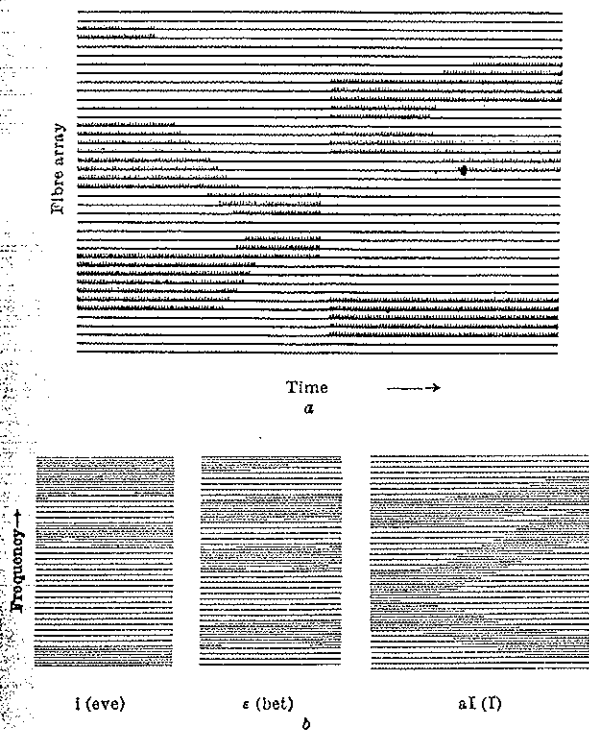


Fig. 5. a, Distribution of pulse activity in the array of fibres in the central auditory pathway, and the way this may vary with time as the components of the stimulus vary in frequency and intensity. b, Distribution of speech energy in different frequency bands for two vowels and a diphthong (after Potter, Kopp and Green²⁰).

It is instructive to compare this with the time course of the formant frequencies of speech sounds (Fig. 5b). The evidence of Fig. 5a also suggests why discrimination between two successively presented tones, which would only require detection of the fact that the edge has moved, is so much more readily carried out than the absolute identification of frequency which would require memory identification of specific channels.

So far, we have considered only steady states and the relations between them. Identifications of frequency changes, however, are also important. For example, the feature which appears to distinguish between the consonants b, d, and g is whether the second formant frequency rises towards its final value, remains constant, or falls towards the final value respectively¹⁷. If we examine

neurons in the auditory part of the cerebral cortex, we find that there are such units which are specifically responsive to frequency changes¹⁸. A given neurone will fire only when presented with a falling frequency, while another will fire only when the frequency is rising (Fig. 6). The precise value of the frequency is largely unimportant provided there is a sufficient excursion in the correct direction. Provided the change is not too slow, the rate of change is not very critical.




Fig. 6. a, Response of a neurone in the auditory cortex to a rising frequency. The signal bar represents a steady tone of 12 kc/s which then rose linearly at the point shown to a new steady value of 13.2 kc/s. The rise time was 50 msec. In b the frequency returned from 13.2 kc/s to 12 kc/s in the same way; this failed to stimulate the neurone (after Whitfield and Evans¹⁸).

It seems likely that the qualities of pitch and intensity are distinguished by levels below the cortex, as discriminations of this type can be made, for example, by cats in the absence of the relevant cortical areas. Temporal patterns of sound, in which the only differences are those of the sequence of the tones and not their identity, cannot, however, be distinguished in these circumstances. There is some evidence of a similar situation in man.

Transient Phenomena and Sound Localization

So far we have been considering either steady state phenomena, or signals involving only fairly slow and limited changes of a single frequency. However, the ear has also to operate on a wide range of transient signals. One of the defects of the resonance theory was its failure to deal with this aspect of cochlear behaviour and its concentration on steady state frequency analysis. It is possible theoretically to produce for any waveform a pure frequency analysis, a pure time analysis, or any intermediate compromise. The mechanics of the cochlea do indeed represent such a compromise between the needs of frequency identification and the necessity for the temporal location of wave fronts. The latter are, of course, of particular importance to the organism in the location of sound sources.

If a short impulse is applied to the ear, a 'travelling bulge' is produced on the basilar membrane which propagates from the base to the apex of the cochlea. This wave has a velocity of the order of hundreds of metres per second near the base and slows to only three or four metres per second near the apex. Experiment shows that as a result of the passage of this bulge a given auditory nerve fibre discharges one or more impulses with a pulse interval corresponding to the steady-state frequency for which that fibre would be maximally excited¹⁹. There is thus a close link between the temporal pulse pattern produced by an impulse and that produced by a low frequency tone, and it would appear that the same mechanism could well process both. The 'click' has a marked superiority over the tone, however, in that it excites temporarily identifiable pulses in fibres at the basal end of the cochlea, and the steep wave front means that these initial pulses will be especially accurately timed.

For localization of a sound source, the relative time of arrival of the sound at the two ears and the relative timing of the corresponding nerve impulses could provide some of the necessary data. In criticizing the telephone and related theories of frequency analysis, we noted the way in which the relation between nerve impulse position and stimulus envelope, fairly good in the auditory nerve, became progressively degraded as more and more synapses were introduced along the pathway. We argued that this potential source of information could only be used if it

were transformed at an early stage into some other code. In the process of sound localization such a transformation does take place. There appears to be a special structure for this purpose—the medial (or accessory) olivary nucleus, which is situated in the brainstem at about the level at which the auditory nerves enter it. The nucleus is a plate like structure which has a sheet of nerve cells each innervated on one face by fibres of the auditory nerve of its own side, and on the opposite face by corresponding fibres of the auditory nerve of the opposite side. Hall²⁰, using click stimuli, has shown that the probability that one of these cells will fire depends on the relative time of arrival of the stimulus at the two ears. Thus, if the stimulus to the ipsilateral ear precedes that to the contralateral ear by 500 μ sec the probability of the neurone firing might be, say, 0.75. If the ipsilateral stimulus lags by 500 μ sec the corresponding probability would be only 0.45. Relative intensity of the stimuli has a similar effect; if the ipsilateral click is more intense than the contralateral one, then the probability of firing is increased and vice versa. Artificial manipulation of the timing and intensity of dichotic clicks enables us to demonstrate the phenomenon of time intensity trading²¹ in the apparent position of a source, but for a real single source the two effects reinforce one another, as the signal will be both earlier and louder at the nearer ear. Because this situation involves a higher probability of a particular cell firing, it will mean that more fibres leaving this nucleus will be activated; the code has been transformed from one in which the information is contained in the relative timing of impulses to one in terms of the number of active fibres. The end result of the transformation thus has the same form, and the same resistance to degradation, as the frequency code discussed earlier. Of course, the number of active fibres on a given side will be a function of the absolute intensity of the signal, as well as of its relative intensity, and to this extent there is an ambiguity. The relative number of active fibres, however, on the two sides is to a first approximation independent of the absolute intensity and thus could give a measure of lateralization. Although such a model has been proposed²², the anatomical site of a comparator has not yet been located. Nevertheless, at the cortex there are cells which, while they are readily activated by quite weak sounds anywhere in the quadrant from dead ahead to a position opposite one ear, are quite unaffected by much stronger signals originating in the corresponding quadrant on the opposite side of the head.

Centrifugal Control

Finally, let me turn to the question of the control of the auditory input. Everyone is familiar with the phenomenon of "not hearing the sitting room clock tick until it stops". Continuously present or irrelevant signals tend to be suppressed at the expense of more immediately important ones. This control takes place quite early in the system. Thus the electrical response in the cochlear nucleus to a sound signal has been shown to disappear if the signal is repeatedly presented over a long period²³. Distraction of attention by the presentation of an "interesting" visual stimulus may have a similar effect²⁴.

These 'gating' operations, which selectively control the input, appear to be mediated by centrifugal pathways which parallel the ascending auditory system all the way from the cortex to the periphery. Even right out at the cochlea it is possible to inhibit activity in individual groups of auditory nerve fibres by stimulation of the centrifugal olivo-cochlear bundle²⁵. Farther up the system, within the brainstem, the possibilities become more varied. Stimulation of one set of descending fibres will shut off the response to sound of a particular neurone in the cochlear nucleus²⁶; on the other hand, stimulation of a different set of these fibres increases the sensitivity of cochlear nucleus cells, and may lower the sound threshold of a particular cell by as much as 15 dB²⁷. It is clear that

these fibres control the sensory throughput at all stages. It is less clear if they simply control whether or not information gets through at all, or if they actually determine the route it takes.

Until quite recently, views on auditory localization tended towards the view that somewhere in the system there should be an area or set of structures uniquely activated by a discriminable stimulus. From this locus activity would presumably 'fan out' to produce a complex widespread response on the motor side. A diagrammatic representation might be made from two pyramids with a common apex. As we have seen, experiment simply does not support this concept of sensory organization. Just why such an arrangement should have been thought necessary is unclear. Among spinal reflexes, the classical 'scratch reflex' involves interconnection of a complex sensory pattern with a complex motor pattern, yet no one has ever postulated a 'scratch' unit or claimed that there were unique pathways leading to it. The 'pyramid' theory, if it were true, would of course have simplified our concept of the discrimination mechanism, as we should only require a gate to operate at a single point. However, the existence of widespread centrifugal connexions means that not only is diffuse gating a practical means of controlling a 'learned' discrimination, but that this gating could take place at almost any neural level. It is well established, for instance, that direct connexions exist between auditory nuclei and motor pathways at the upper medullary level²⁸. These are generally described as 'auditory reflex' pathways, but there seems no reason why they should not be 'conditional' pathways whose activity or quiescence is controlled by the centrifugal system.

In summary, then, let me reiterate two points. There is no evidence anywhere in the auditory system for groups of neurones responding uniquely to particular discriminable stimuli, and multi-channel correspondences probably exist between the sensory system and effector systems, equivalent to those at the spinal level. The elaborateness of the sensory pathway seems to lie in its connectivity; there is no evidence to date for the existence of elaborate code patterns in individual fibres.

¹ du Verney, J. G., *Traité de l'Organe de l'Ouïe* (Paris, 1683).

² Haller, V. von, *Primae Lineae Physiologiae* (Göttingen, 1751).

³ Helmholtz, H., *Gesellsch. Deutsch. Naturf. Aerzte. Amtl. Ber.*, 34, 157 (1859).

⁴ Wever, E. G., *Theory of Hearing* (Wiley, New York, 1949).

⁵ Békésy, G. von, in *Experiments in Hearing* (McGraw-Hill, New York, 1960).

⁶ Tasaki, I., *J. Neurophysiol.*, 17, 97 (1954).

⁷ Schuknecht, H. F., *Neural Mechanisms of the Auditory and Vestibular Systems*, ch. 6 (Thomas, Springfield, 1960).

⁸ Rutherford, W., *J. Anat. Physiol.*, 21, 166 (1886).

⁹ Stevens, S. S., and Davis, H., *Hearing, Its Psychology and Physiology* (Wiley, New York, 1938).

¹⁰ Licklider, J. C. R., and Pollack, I., *J. Acoust. Soc. Amer.*, 20, 42 (1948).

¹¹ Thomas, I. B., *Tech. Rept. 10, Biol. Computer Lab.* (Univ. Illinois, 1966).

¹² Hlilaf, S., and Whitfield, I. C., *J. Physiol.*, 122, 158 (1953).

¹³ Rose, J. E., Greenwood, D. D., Goldberg, J. M., and Hind, J. E., *J. Neurophysiol.*, 26, 294 (1963).

¹⁴ Allanson, J. T., and Whitfield, I. C., *Third London Symposium on Information Theory*, 269 (Butterworth, London).

¹⁵ Whitfield, I. C., *J. Physiol.*, 128, 15P (1955).

¹⁶ Hartline, H. K., *Rev. Mod. Physics*, 31, 515 (1959).

¹⁷ Delattre, P. C., Liberman, A. M., and Cooper, F. S., *J. Acoust. Soc. Amer.*, 27, 789 (1955).

¹⁸ Whitfield, I. C., and Evans, E. F., *J. Neurophysiol.*, 28, 655 (1965).

¹⁹ Kiang, N. S., *M.I.T. Research Monograph No. 35* (M.I.T. Press, Cambridge, 1965).

²⁰ Hall, J. L., *J. Acoust. Soc. Amer.*, 37, 814 (1965).

²¹ Shaxby, J. H., and Gage, F. H., *Med. Res. Council Spec. Rept.*, 166, 1 (1932).
David, E. E., Guttman, N., and van Bergeijk, W. A., *J. Acoust. Soc. Amer.*, 31, 774 (1959).

²² van Bergeijk, W. A., *J. Acoust. Soc. Amer.*, 34, 1431 (1962).

²³ Hernández-Péon, R., and Scherrer, H., *Fed. Proc.*, 14, 71 (1955).

²⁴ Hernández-Péon, R., Scherrer, H., and Jouvet, M., *Science*, 123, 331 (1956).

²⁵ Rex, J., *Acta Physiol. Scand.*, 55, suppl. 139 (1962).

²⁶ Combs, S. D., and Whitfield, I. C., *J. Physiol.*, 188, 34P (1967).

²⁷ Whitfield, I. C., *Bionics Symposium* (in the press) (1967).

²⁸ Rasmussen, G. L., *J. Comp. Neurol.*, 84, 141 (1946).

²⁹ Whitfield, I. C., *Brit. Med. Bull.*, 12, 105 (1956).

³⁰ Potter, R. K., Kopp, G. A., and Green, H. C., *Visible Speech* (van Nostrand, New York, 1947).