NEUROLOGICAL MODELS AND INTEGRATED PROCESSES *+ [158]

Warren McCulloch, Michael Arbib, Jack Cowan

Introduction

We of the Neurophysiology Group of the Research Laboratory of Electronics, M.I.T., are, and have been, interested primarily in constructing models of the working of the nervous system, not models of the way it came to be what it is at the moment.

In this paper we wish to survey what we have accomplished in the last 4 years. We shall not go into technical details, for formal definitions, theorems and proofs of all our results have already been published elsewhere, or will be published in the near future.

Our work has proceeded from the grossly simplified neural model of McCulloch and Pitts, 1943,[1] which ignored the effects of disturbances both in the functioning of individual neurons, and in the organization of the neurons into networks. Our recent efforts have been directed to constructing more sophisticated models, albeit still grossly abstracted from neurology, which are designed to combat such noise.

We are chiefly concerned with modelling the functional organization of the nervous system. Hence, its anatomy is significant for us in two senses—first, it puts structural limitations upon hypotheses of circuit actions, and second, it defines the sites where we should excite and whence we should record activity.

For reasons of economy we have adhered strictly to an electrical hypothesis of central inhibition and excitation.[2] This requires that the electrical properties of transmission and the geometry of the structures determine the output as a specific function of the input to each tissue.

Perhaps the simplest example of this unique correspondence of structure to function is located in the superior olive whose large cells receive signals from both ears. The anatomy of these cells, and of the synapsis upon them, is such that signals from either ear alone excite them; whereas contemporaneous signals from the two ears cancel each other, leaving the cell unaffected. By these cells we detect the direction whence a sound comes and hear a signal in one ear despite much noise common to both.

A second clear example is to be found in the cerebellum which evolved from an arch of large cells, each receiving signals from both vestibules, conveying the accelerations of the head. The input has come to include signals from almost all parts of the body and from the rest of the brain. They arrive on a host of small cells whose fine and uniform axons divide to run transversely through the branches of the giant cells, as wires through parallel rows of regularly spaced telephone poles, some quarter of a million wires contacting each pole in its row. The cerebellum serves as an interval clock in all our ballistic acts, such as fiddling or writing, in which accelerations and decelerations must be precisely equated and timed—and that at intervals too brief for reflexes to mediate the process. But it also serves for autocorrelation of signals of various modalities to combat noise. The modality varies according to the bird, beast or fish, often as an aid in flight, balance or navigation. These specializations are rather a function of the particular inputs than of the structure, which is curiously uniform in detail of synapsis.

Our reason for considering the superior olive and the cerebellum is that the function of the tissue does not inhere in any particular neuron but is distributed over all the neurons. Anatomically, this is recorded in the multiplication of similar components and in the multiplication of overlapping inputs. Physiologically, it is obvious in the time domain, for both tissues operate, one in detecting and the other in timing, with a precision of l-2μsec; whereas the components, either in detecting coincidences or in emitting impulses, have a standard deviation greater than 1/3 msec. Pathologically, it follows from the negligible loss of precision, despite the scattered death of one-tenth of the cells. *We emphasize that both structures exhibit a stability of performance greater than that of their components even though each component is computing essentially the same function, albeit of a somewhat different set of inputs.*

Our second point is better exemplified in the frog's eye.[3] The adequate stimulus for his rods and cones are light of appropriate wavelengths, but the adequate stimulus for any bipolar cell is a certain figure in space and time of that light, which depends in detail upon the junctions of the receptors to each bipolar. The frog has four varieties of ganglion cells differing in size and in pattern of branches and, of course, in patterns of synapsis with bipolars. These are such that it is possible to assign to each kind, which of four sorts of visual properties it will respond to. Each yields one quality of one spot, be it small or large. For perception, such a cognizance by adjective is insufficient without cognizance by relation. In the frog the latter is supplied by the mapping of the four qualities in four layers in the superior colliculus, preserving in each the topological properties of the retina. These four maps are in register, and when the frog's optic stalks are cut, they grow back so as to reconstitute the four maps in

register. Such a self-reorganization requires, for a million ganglion cells, at most, 20 bits each of selective information—far less than 0.1 per cent of the information supposedly stored in the structure of each of these cells. This should be no shock to the embryologist who knows that in certain lowly forms of life, a transplanted midbrain can generate a whole forebrain. It does point to the question of how much information one had better put into a structure that is to receive information from its world and process its signals so as to know that world well enough to survive.

There is reason to believe that the precise location and sequence of connections to dendrites and cell bodies, as well as the interaction of afferents determined by their juxtapositions, specify what function each cell is to compute at a particular threshold when the signals it receives are at specific strengths and specifically timed. As every experimental neurophysiologist knows to his sorrow, there is no reason to believe that any one of these is exactly determined by our genes, by adaptation, or by learning.

Let us begin with the detailed synapse of our 10^{10} neurons which can scarcely be specified by our genes under fixed proper environmental conditions. Neuroanatomists generally believe that an appreciable number of axons normally grow astray. Estimates of 1 per cent, or more, are commonly made for various structures. Adjacent similar columns of cells in the cerebral cortex are never exactly alike in their anatomical pictures, and it seems reasonable to believe that, while their synapsis is generally specified, the details are left largely to chance. Even if synapsis were specified in detail, the random death of at least one neuron a minute would soon disorder the details. Myelinatlon of finer axons, at least in the outer layers of the cerebral cortex, continues to at least the 50th year of life. Moreover, the brain exhibits visible pulsations with respiration, with heartbeat and at a higher frequency, recently attributed to rhythmic contractions of glia. Certainly a living brain has no fixed or rigid geometry.

When we turn to other parameters affecting nervous activity, we note that many general chemical changes alter all components similarly. For example, homeostatic mechanisms tend to keep the brain at a pH of approximately 7.2. If it rises to 7.4, the threshold of neurons, soma and axon, falls to approximately SO per cent of the normal value, at which alkalosis, one begins to have carpopedal spasm. When the pH falls to 7.0 the thresholds rise 100 per cent. These are little more than the changes one can induce safely by hyperventilation and by holding one's breath, and neither of them prevents a diver from performing complicated tasks. Under surgical anaesthesia with ether, the pH of the brain is approximately 6.9, but the respiration continues automatically. These changes that are due to pH depend primarily on the alteration of ionization of calcium, but in buffered preparations of nerves, Lorente de No has shown a specific effect of carbon dioxide approximately proportional to the logarithm of its concentration. It raises both the rate of repolarization after the transmission of an impulse and the voltage, thereby increasing the strength of the impulse and decreasing its rate of propagation.

Finally, thresholds and strengths of impulses are sensitive to temperature. Yet human brains have been known to work, albeit not too well, at 42 °C,[4] not far beneath the lethal temperature. Below 26°C, mammalian nervous tissue becomes unexcitable—but at several degrees above this temperature patients, chilled for cardiac surgery, can still think. We have been particularly interested in constructing models of nervous circuits whose input-output function remains constant under these common shifts of threshold. Von Neumann called them "logically stable" circuits. They are mentioned in Agatha Tyche[5] (We shall use "component" and "line" to designate our formal analogues of "neuron" and "axon.") Manuel Blum has been able to design them for components with any number of inputs, stable between those limits enforced by the output component escaping control because it fires for every input, or because it can never be fired. He will publish his findings soon, but we should note in passing that the construction of formal nets that are logically stable over the whole usable range requires that inhibiting interaction of afferents which enables a component to compute any logical function of its inputs, and to have these functions follow each other in any possible, prescribed sequence as its threshold shifts. That same interaction removes from the real neuron the restriction to compute only those functions that are computable by threshold logic.

It has long been known that rapidly repetitive activity of nerve cells and of their axons affects their recovery and, consequently, both their thresholds and the strength of their impulses. Moreover, during recovery, their voltage gradients are known to affect other neurons In their vicinity. These effects are large and easily demonstrated when many components are fired in unison. They must occur normally, to some extent, and may well account for much of the large fluctuation of threshold which we detect in them; but whether these jitters, say some 10 per cent of threshold, are used in signalling or are only noise is not yet decidable.

Be that as it may, one has to admit the possibility of many local random changes for reasons like those responsible for general changes, and for many other little accidents. Finally, since every trigger point is a small area of high specific resistance and must operate at body temperature, the threshold must jitter. The best measurements of this type of biological noise are those of Dr. Verveen, made on excised nerve under carefully controlled conditions. These are his conclusions—and we quote him verbatim.

"A nerve fiber responds to the application of an electrical stimulus with the production of an action potential in a fraction of all trials. The relationship between this probability of response and stimulus intensity approximates the Gaussian distribution function and is characterized by its coefficient of variation (the so-called Relative Spread (RS)). The value of this coefficient was found to depend on the diameter of the axon. In large axons the RS is 'small' and In small axons its value is 'large.' It varies from 0.5% in 250 µ thick giant axons of sepia to 1% in 4µ thick axons of frogs. The actual range of the threshold region as compared to the value of the threshold is about 5 times the value of this parameter. This experimental relation bears a close

He would have you note that this jitter, being of high frequency, effectively lowers the threshold, although we are wont to think of it only as thermal noise. One can imagine other uses for it in the feltwork of the brain.

In any case, any adequate neurological model must be designed to function properly despite local random variations. In his Master's Thesis, Manuel Blum[65] has shown how to design such nets with one rank of any number of components each with inputs from the same number of sources, and all of the first rank playing upon a single output component. He presented some of his conclusions to the first Symposium on Bionics.[7] Briefly, for 2 inputs each, if every component may misbehave for even a single configuration of its inputs, there is no error-free behaviour except in computing tautology or contradiction. Such behaviour begins with 3 inputs and increases rapidly with the number of inputs (and with the number of components in the first rank) so that with a hundred, those of the first rank need only behave properly to 8 per cent of their input configurations, and the output neuron, to only 2 per cent of its input configurations. The important thing to note here is that it does not matter with what frequency a local error is committed, or whether it depends upon threshold, strength of stimulation of details of synapse—no error occurs in the output. The one requirement is this: that, for the few configurations of inputs to which each component is required to behave properly, it does so. Hence, for these configurations, the input to each component still determines its output.

No similar error-free behaviour is possible when components escape completely from their inputs. Here we should distinguish two kinds of failure—the first occurs in the case of dying neurons which often in disease emit long trains of impulses where there should be none and when dead, emit none when there should be some; the second kind of misbehaviour is the scattered distribution in time and place of impulses arising in axons (lines) or their failure to propagate for no good reason.

To cope with these difficulties we first tried von Neumann's[8] scheme of replacing single lines by bundles each line of which carried the same signal, but instead of going to only one component, each line made connections with every component in a "rank," each of which computed the same function, i.e., a "rank" of n complex components computed the same function that only one ideal component would do. Such a network comprising complex components is said to be *redundant*, and it is this redundancy of complex components which permits reliable computation to be performed by unreliable components. The result of using such complex redundant networks was that a marked improvement was obtained in reliability and amount of redundancy required, compared with von Neumann's scheme (Verbeek et al[9, 10]). This solution, promising though it was, was not satisfactory in that extremely reliable computation was obtained only by using highly redundant networks, similarly to von Neumann's scheme.

Shannon's famous noisy coding theorem[11] indicated that, if It could be applied, such reliability could be obtained by using networks with much lower degrees of redundancy, given arbitrarily complex components with small probabilities of malfunction. We have recently demonstrated (Cowan and Winograd[12]) the validity of this theorem under certain assumptions for computing networks, so that error-correcting codes (Shannon[11]) may now be used to combat failures and malfunctions of computing components.

Briefly, the solution consists of replacing not one but many simple components, computing different functions, by a rank of complex components each member of which is computing, in general, a *different* function of its inputs. The network composed of such assemblies of components may be made arbitrarily reliable, apart from errors in the final outputs which cannot, of course, be corrected, by ensuring that each rank corrects for errors in the previous rank. Again, high reliability of computation is achieved by redundant networks of highly complex components, but the redundancies may be orders of magnitude smaller than those required by von Neumann and Verbeek et al. Such networks comprising many complex components each computing different functions have been termed *anastomotic*.

Provided only that such complex components are not more noisy than simpler ones, Shannon's theorem may be shown to be valid, i.e., minimum redundancies, determined only by the component noise, may be used to achieve arbitrary reliable computation by networks of unreliable components. Moreover, such networks need not be precisely connected, i.e., a limited amount of perturbation of connection can be tolerated. If component errors do increase with complexity, then higher redundancies are required, but the amount may be kept to a minimum.

The application of such theories to neurological models thus furnishes some specific questions. Are larger neurons in the nervous system more stable than smaller ones? Are more complex anastomotic nets more reliable than less richly connected ones? It seems reasonable that large neurons are in fact "quieter" and more stable than smaller ones—the richest connections are to the largest neurons whose thresholds should be the most stable. For these large cells in the cerebellum, the cerebrum, and the spinal cord, an estimate of something a little less than a quarter of a million connections is certainly reasonable. Similarly, if a general estimate of the fan-out of inputs be desired, one has only to remember that a man has a few million inputs to his nervous system, and that a single perception usually requires say one-tenth of a second, and, hence, not more than, say, 100 neurons in depth. Similarly, one has a few million outputs to muscles and glands on which signals must again converge.

Certainly less than half of our neurons are concerned only with feedback toward the receiving sides, consequently the fannings are on the average at least of the order of 1 to 1000 inward and 1000 to 1 outward. Such is the general picture which compares the few cells necessary for a computation with the many available for coding.

The cortex of the cerebellum is certainly a special-purpose computor built to do the same thing throughout our lives. Per contra, the cortex of our cerebrum is born immature in structure, as well as in function, and we know that its normal growth depends upon its input. In those born blind its visual area at maturity is thinner than in normal man by about a third. At birth, cortical neurons are still sprouting branches and it is thought that, whether it be a matter of maturation, use or learning, their number decreases, the larger branches growing at the expense of the smaller. This has been demonstrated by Gairns on sympathetic ganglia in which it is clear how the change converts a diffuse response into a highly specific one. Presumably the same holds for the cerebrum, but in ways too complicated for us to recognize easily. Functionally, we see its parallel in the initial generality of a conditioned reflex differentiating with enforced discrimination. In such an adaptable structure only generalities as to the distribution of kinds of neurons, and of the directions of their axons, need be specified, the initial branching being left largely to chance, and the final connections determined by subsequent events, the exact picture depending upon the temporal sequence of those events. The resulting picture of the visual cortex would then be as Sholl described it, i.e., only statistically regular. Though its inputs do map topologically the retina and hence the visual field, its projections upon subsequent cortical structure appear to lack all such correlation.

Anatomists relating behaviour to brains, have found that creatures capable of more complex, flexible and reliable activities have a larger portion of the volume of the cortex devoted to the ramification of connections of its diversified components. With a few possible exceptions, including the dolphin, the ratio of cell bodies to total volume is smallest in humans. Provided that the diversity of connections be great enough, there will be some which will grow with maturation and use. There is thus in a good cerebral cortex the adequate variety for reliable computation despite noise of every kind.

Although our concern has been with the way in which messages are preserved and compounded in systems that have organized themselves, rather than in how they evolved, certain general features seem to emerge which find their counterparts in the process that produced them—perhaps because at every step in that process the system had to survive, so that the present creature carries a message, or specification, that evolved in the past. When we speak of a system we have in mind something whose parts stick together long enough for us to recognize it. For this, first, there must be the requisite variety of parts; and second, these parts must be sufficiently strongly tied to each other to stick, for all natural things are products of an entropic process (i.e., one obeying the 2nd law of thermodynamics), even as seemingly improbable a creature as a warthog.

Next, one notices that structures that have regularity of order and form, be they crystals or bacteria, often promote the reduplication of that form in the environment where they arise, this autocatalysis yielding the ancient law that like begets like.

In simpler systems only accidents can produce changes from among which the world makes its selection, but in higher forms, there is a planned shuffling of messages conveyed by chromosomes that produces a variety of types, thereby increasing the viability of the kind. This, the hereditary net, has its counterpart in the anastomotic nets of our blood vessels, and in the nervous nets of the central nervous system.

To be a self-organizing system the system must make organs, or organelles, to perform specific functions. For this, each organ, and its components, must sacrifice some of their former versatility, just as a general Turing machine must, if it is to become a special one. We may instruct the machine with a program on a tape; whereas, while nature may do the same with its genetic determinants by masking parts of them, what appears is a special-purpose computer with the information stored in its structure.

These organs nature builds to select its inputs so that they are useful to the system, and to perform those routine acts by which its effectors cope with the world. This differentiation is a one-way path from the totipotent zygote to the neuron that cannot reproduce and the erythrocyte that has no nucleus. In the nervous system we find it wherever complexity of adaptable functions must be sacrificed for precision and celerity. Evolution has not multiplied them beyond necessity, nor varied wantonly their regular construction. If, in perception or execution, one has in mind any practical and interesting job for which one would initiate a self-organizing system, he had better program, or solder in, these necessary structures. He who wants a sweetheart in the Spring would not be wise to wait for an amoeba to evolve her—and there is no use starting with a hundred amoebae instead of one.

Yet it is well to remember that our artefacts are an order imposed by us on things not so ordered naturally. As MacKay says—we lack nature's glue. Imposed order is the beginning of an entropic process that destroys the artefact. Natural order is the end of an entropic process—even in that improbable example in which the natural object is a sweetheart.

Footnotes

References

1. McCulloch, W. S. and Pitts, W. H., "A Logical Calculus of the Ideas Immanent in Nervous Activity." Bull. Math. Biophys. 5, 115-133 (1943).

2. McCulloch, W. S., Lettvin, J. Y., Pitts, W. H., and Dell, P. C., "An Electrical Hypothesis of Central Inhibition and Facilitation." Chapter V in Patterns of Organization in the Central Nervous System. Res. Publ. Assoc. Res. Nerv. Ment. Dis. XXX, pp. 87-97, December 15-16, 1950.

3. Lettvin, J. Y., Maturana, H. R., Pitts, W. H., and McCulloch, W. S., "Two Remarks on the Visual System of the Frog" in Sensory Communication (The M.I.T. Press, Cambridge, Mass., and John Wiley and Sons, Inc., New York, 1961), pp. 757-776.

4. Simons, Lt. Col. David G., Man High. (Doubleday, New York, 1960).

5. McCulloch, W. S., "Agathe Tyche, Of Nervous Nets... The Lucky Reckoners" in Mechanisation of Thought Processes. (H. M. Stationery Office, London, 1959), pp. 611-625.

6. Blum, M., "Reliability of Biological Computers." S. M. Thesis, Department of Electrical Engineering, M.I.T., June, 1961.

7. Blum, M., "Properties of a Neuron with Many Inputs." Proc. Bionics Symposium, pp. 55-82.

8. von Neumann, J., "Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components." Annals of Mathematics Studies No. 34 (Princeton University Press, Princeton, N. J., 1956), pp. 43-98.

9. Verbeek, L. A. M., "Reliable Computation with Unreliable Circuitry." Proc. Bionics Symposium, pp. 83-92.

10. Cowan, J. D., "Toward a Proper Logic for Computation in the Presence of Noise." Proc. Bionics Symposium, pp. 93-152.

11. Shannon, C. E. and Weaver, W., A Mathematical Theory of Communication (University of Illinois Press, Urbana, 1941).

12. Cowan, J. D. and Winograd, S., "Functionally Stable Automata" (Submitted to Phil. Trans. Roy. Soc. (London)).

13. Proceedings of the Bionics Symposium, Dayton, Ohio, September 13-15, 1960. WADD Technical Report 60-600, March 1961.