


SPECIAL TELEVISION NUMBER

HUGO GERNSBACK, Editor

RADIO — ELECTRONICS

formerly

**RADIO
CRAFT**



THE SPIRIT OF TELEVISION

MAR

1949

50¢

U. S. and
CANADA

LATEST IN RADIO — ELECTRONICS — TELEVISION

The Electronic Brain

By W. R. ASHBY, M.A., M.D.*

TWENTY years ago the idea of building a brain would have been considered fantastic. Mind and matter had been carefully separated by the philosophers who were mostly convinced that any non-living connection was impossible. No mere machine, they said, could produce the remarkable features of the brain. In a sense, of course, they were right. When they thought of a machine they imagined objects like a wheelbarrow, a typewriter, or a steam-engine. They had observed that such machines if controlled, like a typewriter being tapped, were inflexible in action, and if uncontrolled, like a motor car without a driver, were apt to destroy themselves.

But nowadays the word "machine" has a much richer meaning, the position having been transformed by the invention of the electron tube. This device has two main properties: it allows power to be injected freely into a machine, causing high activity, and it provides a means by which one part of a machine can affect the behavior of another part with little back-action. At last those who would build a brain have something comparable in functioning powers with the nerve cell.

The nature of a brain

But even if we are given an abundance of highly active and sensitive devices like nerve cells or tubes we have yet to assemble them into something that makes sense. And what does "make sense" mean in the brain or in a machine? Here wide differences of opinion occur. To some, the critical test of whether a machine is or is not a "brain" would be whether it can or cannot "think." But to the biologist the brain is not a thinking machine, it is an *acting* machine; it gets information and then it does something about it. Like every other organ in the animal body, it is a means to survival.

This last property decides its fundamental mode of construction; it must have certain permanent goals—the essential conditions for its existence—and it must be able to attain them in a variety of circumstances. If one path to the goals is blocked it must find another. If the circumstances change, it must readjust its methods. The brain of an insect has available a few perfected inborn patterns of behavior. It will try them in turn, succeeding if the circumstances are of a standard type.

The brain of a mammal is of more interest to us, for it has a diffuse ability to puzzle out some sort of adaptation to an indefinitely large variety of circumstances. Man is himself the outstanding example of the potentialities inherent in this subtle mechanism.

The ability of the brain to look after itself by correcting all deviations from

Here, for the first time in any American magazine, we present to our readers the complete account of the new epoch-making and thought-provoking electronic brain, developed in England, and featured recently in news dispatches all over the world.

a central, optimal state, and particularly its ability to do so by a variety of methods, being flexible about the route but unchanging in its aim, was usually regarded as quite beyond the powers of any machine, but it has been known since 1940^{1,2} that machines of the more dynamic type can do this quite easily. All that is needed for this goal-seeking flexibility is that the system should have negative feedback. When there is no feedback, as when a door-bell button is pressed and the bell rings, there is neither sense nor nonsense—it does just what its present state of repair or disrepair enforces. But when a radar-controlled anti-aircraft gun receives impulses both from the target plane and from its own shells, and is affected by the distance between the two so that it tends to make the distance between the two zero, then such a system has negative feedback and is "goal-seeking." The important point here is that *the property of being "goal-seeking" is not that of life or mind but of negative feedback*, and any machine, however inanimate, which has negative feedback will show this feature.

Self-organizing machines

But this does not complete the requirements. Thus, if the gun-radar-plane system had *positive* feedback it would tend to make the distance between shell and plane a maximum and

would therefore seem to be trying to get its shells as far away from the plane as possible. Clearly, the introduction into a system of feedback in general does not solve the problem; for if without feedback the gun will aim anywhere, yet even with the feedback it may either seek the target or it may positively avoid it. What is to ensure that the feedback has the correct sign?

In the gun-radar-plane system the problem is easily settled: the designer carefully arranged the construction so that the feedback was negative. In the brain of an insect, all variations born with wrong feedbacks were eliminated by natural selection ages ago. But in the higher animals the position is different. Large numbers of the feedbacks are left at first undecided, since it is experience and not the inborn (genetic) characters which are to determine the feedbacks. Thus, a cat may have to learn to go towards red meat (negative feedback), but to go away from red embers (positive feedback).

That a kitten's initial feedbacks are rather chaotic is shown by the way in which it may shrink away from a saucer of milk and then run towards a red-hot fire. Yet we know from experience that day by day the kitten's feedbacks change, always improving, and tending to those values, positive and negative, which ensure the animal's survival. The problem of the mammalian brain, then, is that as a machine *it has to work out an essential part of its own wiring*.

The homeostat

Such ability to learn and to adapt by internal re-organization was regarded as a great mystery, but the principles are now better understood.³ To demonstrate them and to show that these principles do, in fact, produce such behavior, a machine has been constructed and has recently been demonstrated.⁴

The homeostat consists of units, four of which are shown in Fig. 1. Each carries on top a suspended magnet, shown in Fig. 2, and the behavior of these four magnets provides the focus of interest.

Each magnet (M in Fig. 3) is affected by currents in the four coils around it, the currents coming partly from the other units (A, B, C) and partly as a self-feedback (D). (The apparently single coil of Fig. 2 is composed of the four coils of Fig. 3). In front of each magnet is a trough of water with electrodes at each end at -2 v and -15 v respectively. The magnet

* Department of Research, Barnwood House, Gloucester, England.




Fig. 1—The homeostat, with its four units, each one of which reacts on all the others.




Fig 2—Quadruple coil ABCD encircles magnet M which is suspended by the needle pivot. The suspending wire extends forward on its end into the water in the semicircular plastic trough which has electrodes at each end. Potential for the grid is taken from the pivot socket.

is suspended on a needle pivot by a wire sling which dips into the water and picks up a potential which depends at each moment on the position of the magnet. The potential goes to the grid of a triode and thus controls the d.c. output of the unit. (The resistor E is first adjusted so that when the magnet is central the unit has zero output). This output goes to the other units in series where it becomes one of their inputs.

This arrangement sets all four units into action and reaction on one another. As soon as the system is switched on the magnets are moved by the currents from the other units, but these movements change the currents, which cause fresh movements, and so on.

These actions and reactions can be modified by various constant settings. Thus, the current from, say, unit 4 to unit 2, can be controlled as to its po-

larity of entry into the coil by X (Fig. 3). In addition, the potentiometer P decides what fraction of the input current actually goes through the coil. These controls can be hand set by the upper two rows on the front panels.

When set in some way, the magnets show some definite pattern of behavior, the pattern depending on the pattern of the hand settings. If these latter give a stable arrangement then the four magnets move to the central position where they actively resist any attempt to displace them. If displaced, a coordinated activity brings them back to the center, rather as an animal positively seeks its optimal conditions. Other settings may, however give instability, in which case a "runaway" occurs and the magnets diverge away from the centers. In such cases the feedbacks are producing "vicious circles" which would be driving the animal away from its optimal conditions.

But the feedbacks, instead of being set by hand, can be controlled by similar wirings arranged on a uniselector (V) in each unit. The values chosen for the wirings were deliberately randomized, the actual numerical values being taken from a published table of random numbers.⁵ When controlled by the uniselectors, the pattern of feedbacks depends at any moment on the values provided by the uniselectors at that time. Twenty-five positions on each of four uniselectors means that 390,625 combinations of feedback patterns are available.

Finally, in each unit the uniselector moves to a new position when and only when the output current of that unit exceeds the value sufficient to close the relay (F), the latter energizing the coils (G) of the uniselector.

When the control is diverted by the switches S-S so that not the hand controls but the uniselectors determine the settings, then a new feature emerges in the behavior of the system. As before, the units start acting on one another, but the uniselector settings change whenever the system is unstable, *i.e.*, whenever the magnets diverge far from the central position. In other words the machine starts to hunt for a combination of uniselector settings giving a stable system, *i.e.*, giving the proper internal feedbacks. When it finds a combination with the right feedbacks it holds that combination and will then demonstrate that it has assembled that feedback system which results in a coordinated maintenance of its variables at optimal values, like a living thing. The important point is that it finds its own arrangement of feedbacks, the designer having merely provided it with plenty of variety.

Not only will it find the appropriate feedback initially, but if we alter the basic conditions in any way it will proceed to re-adapt itself to the new conditions. Thus, we may use hand controls on two of the units, setting them at arbitrary values to represent some "environment" to which the other two units, representing "nervous system,"

must adapt, i.e., find combinations of their two uniselector settings which, in relation to that particular "environment," forms a stable system. When the machine is switched on, it proceeds, as described above, to find such an adaptation. But if now we alter the hand settings, i.e., change the "environment" to which the other two units are adapted, then the machine promptly abandons those uniselector combinations and hunts for new ones which will restore adaptation to the new environment. If now we change the hand-settings again, a new appropriate combination will again be found. And this process can be repeated as often as we please.

But the homeostat will adapt not only to random changes in hand settings but to any change in the dynamic nature of the machine, whether of a type originally intended or not. Here, for instance, are some alterations suggested by my colleagues who have tried to confuse it. After it had found a stable combination we reversed the polarity of the connection of an output to an input; it promptly changed its uniselector settings till it found a new combination of settings which was stable in conjunction with the new conditions. We reversed the polarity of a trough, thereby changing some of its feedbacks; it changed its uniselector settings till it found a new combination of settings stable in conjunction with the new conditions. A magnet was reversed: it readapted to the new condition. Bars were placed across the troughs so that the magnets could swing only to one side: it readapted. We joined two of the magnets together with a light glass fiber so that they had to move together; it readapted. In all cases, whatever conditions were imposed, it rearranged its own wiring through the uniselectors until it developed the proper feedbacks in relations to the new conditions.

Is the homeostat a brain then? Hardly, for it is as yet too larval. But it uses a new principle and can easily be extended to give much more powerful developments. Its chief fault in its present form, with only four units, is that it has little room to accumulate new adaptations, but, if it has to adapt to a new environment, must obliterate its established adaptations to make room for the new. This, of course, is a serious handicap, just as a child would be handicapped at school if it could learn what was two times three only by losing its memory of what was two times two. The difficulty, however, is a minor one and could be overcome by a mere increase in the number of units together with some minor alterations.

The making of a synthetic brain requires now little more than time and labor. But there is one point on which we must be quite clear: a proper synthetic brain must develop its own cleverness—it must not be a mere parrot. No matter how dazzling the performance, we must always ask how much of the performance has been enforced in detail by the designer and how much is contributed by the machine itself.

Let us suppose that two machines have been developed to the point where they can actually play chess. First we consider an electronic computer of the ACE or ENIAC type.^{6,7} Instructions may be fed into it so that it will make only legal moves, but this is insufficient—a random series of legal moves will not win games. The machine may have great powers of analysis, but unless this ends in a demonstrated mate, the analysis

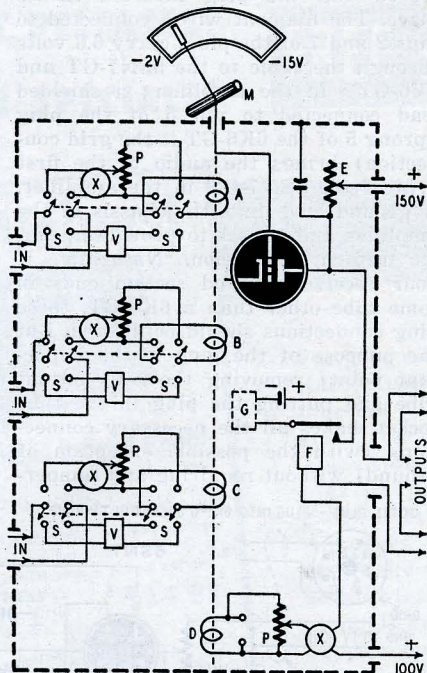


Fig. 3—Schematic diagram of the homeostat.

must stop at a judgment. (I assume that chess, like living, cannot always be analyzed out completely.) If the designer supplies it with criteria for judging whether positions are to be aimed at or avoided, then the criteria must be decided by the designer. This being so, such a machine, if perfect, will produce chess based on a strategy as good as its designer's but no better.

The second feature of such a system is that its thousandth game will be no better than its first.

The third feature is that every part has an exact duty set by the designer, who can say at any moment whether it is or is not working in accordance with his design and instructions. In short, it is a slave-brain.

The other type of machine, the homeostat, is based on quite a different principle. It needs no detailed instructions, only some method by which it is informed of the occurrence of illegal moves and mates. How the machine is to avoid these undesirable informations (feedbacks) is left entirely to the machine to puzzle out for itself. (The adaptations already shown by the homeostat encourage the confidence that with only minor developments the machine will succeed.)

Let us suppose the homeostat perfected and contrast its behavior with that of the first machine. The homeostat would start off like any other player—simply by making more or less

random movements. But the feedback would soon stop it making illegal movements, and it would tend steadily to avoid the moves that lead to a rapid loss of the game. But it must be admitted that its first games would be very bad—as bad, in fact, as the first games of any future world champion. But the homeostat would tend steadily to shed bad moves. Lines of play would be developed or dropped simply according to whether they did or did not lead to a win.

These improvements would be in no way dependent upon the particular details provided by the designer: they would be developed by the machine out of the indiscriminate variety provided, the feedback being the dominating and controlling factor. Consequently, such a machine, if perfect, could eventually play with a subtlety and depth of strategy beyond that of the man who designed it.

The aim of some has been to produce the perfect slave-brain. Though undoubtedly useful for some purposes, yet we must not lose sight of our objective: a synthetic brain should not only play chess, but should eventually beat its own designer.

This prospect is now in view.

And the future?

And what after that? Some facts seem clear even at this distance.

The construction of a machine which would react successfully to situations more complex than can be handled at present by the human brain would transform many of our present difficulties and perplexities. Such a machine might be used, in the distant future, not merely to get a quick answer to a difficult question but to explore regions of intellectual subtlety and complexity at present beyond the human powers. The world's political and economic problems, for instance, seem sometimes to involve complexities beyond even the experts. Such a machine might perhaps be fed with vast tables of statistics, with volumes of scientific facts and other data, so that after a time it might emit as output a vast and intricate set of instructions, rather meaningless to those who had to obey them, yet leading, in fact, to a gradual resolving of the political and economical difficulties by its understanding and use of principles and natural laws which are to us yet obscure.

The advantages of such a machine are obvious. But what of its disadvantages? There are at least two.

Firstly, in its construction, many a detail will have to be fixed at some arbitrary value selected without full knowledge of what features it may impose ultimately on the type of reactions. Are the tubes, for instance, to have an anode voltage near the maximum or much lower? Once made, such a decision will result in an all-pervading tendency in the machine's behavior. One machine, for instance, might try to solve all problems by exploring the possibilities of immediate violent activity,

