BRAINS, INTELLIGENCE, CREATIVITY, AND GENIUS

INTRODUCTION

In our increasingly complex and interconnected world one sometimes hears wistful expressions to the effect that what we need is some superintelligent leader who can cope with it all on our behalf, since our own intelligence seems so inadequate. As Ashby points out in "Design for an Intelligence Amplifier," this is a forlorn hope, since human intelligence is severely limited. In a fruitful analogy he points out that mankind's big advance in mechanical strength did not come from breeding burlier slaves but rather when Watt discovered a way to tap natural sources of power. The steam engine is a power-amplifier in the sense that a small amount of power (used to manipulate the controls, fuel the fire, etc.) is used to control a much larger amount (liberated from the fuel, ultimately). For intelligence-amplification, then, we need some way to use human intelligence, which Ashby argues is basically selection, to control much larger sources of selection.

In the field of artificial intelligence, which was in its infancy when the article was written, the basic process outlined by Ashby is used very commonly. Always the goal of intelligent action, by human or machine, can be viewed abstractly as an appropriate selection from a set: a chess move from all moves legally possible; the answer to a test question from the set of all available answers; a sequence of actions which will accomplish an objective, taken from the set of possible actions, and so on. To generate a set of moves, answers, or actions is rather easy and no sign of great intelligence; to select the right one(s) from the set is the hard part. When we write (or abstractly, select) a computer program to do the selection for us we are obtaining selection-amplification. To the extent that the essence of intelligence is selection, we are obtaining intelligence-amplification, and there is no problem in devising a machine which exhibits more intelligence than its designer.

In this farsighted and penetrating article Ashby also discusses solutions to some of the problematic combinatorial aspects of "selection by equilibrium." All of these are reflected in modern efforts in artificial intelligence.

"Can a Mechanical Chess-Player Outplay its Designer?" From the arguments in the first article in the chapter one can conclude that the answer is 'yes,' and it is also an historical fact that computers do outplay their programmers. (Ashby would have been delighted to hear that as this book was being edited, a computer program defeated the world's champion of backgammon, a substantially difficult intellectual feat.) The question addressed by the paper, however, is broader: can a machine exhibit more "design" than is put into it by its designer? The famous "Argument from Design" for the existence of God is based on the allegedly "self-evident truth" that it could not, so the question is of philosophical interest. Ashby casts the philosophical question in modern cybernetic terms and shows that if the designer contrives the machine in such a way that it can use information from its own experience, then it can eventually show "more design" than that put in by its designer. The true Designer of a machine is thus its explicit designer, together with the environment which by sending it information allows the machine to enhance the original design.

"What is an Intelligent Machine?" In this and other articles Ashby takes to task the notion that the brain is a magical device with supernatural skills, arguing that the comparison between "smart" brains and "stupid" computers is stacked against the machines when the intelligence demanded is matched to the everyday environment for which we have been highly preprogrammed (designed, in the language of the prior article) both genetically and through decades of learning. For intelligence is basically goal-seeking activity and appropriate selection, and for this both brains and machines are subject to the same limits, such as the Law of Requisite Variety. A high intelligence is shown by a device which utilizes information efficiently to achieve a high degree of appropriate selection. Even the so-called "genius" cannot escape the laws of information which govern the process. In this excellent article Ashby gives us an exceptionally clear and crisp discussion of the subject of intelligence, as applied to brains and synthetic machines.

DESIGN FOR AN INTELLIGENCE-AMPLIFIER

SECTION I

1. Introduction

For over a century man has been able to use, for his own advantage, physical power that far transcend those produced by his own muscles. Is it impossible that he should develop machines with "synthetic" intellectual powers that will equally surpass those of his own brain? I hope to show that recent developments have made such machines possible – possible in the sense that their building can start today. Let us then consider the question of building a mechanistic system for the solution of problems that are beyond the human intellect. I hope to show that such a construction is by no means impossible, even though the constructors are themselves quite averagely human.

There is certainly no lack of difficult problems awaiting solution. Mathematics provides plenty, and so does almost every branch of science. It is perhaps in the socail and economic world that such problems occur most noticeably, both in regard to their complexity and to the great issues that depend on them. Success in solving these problems is a matter of some urgency. We have built a civilization beyond our understanding and we are finding that it is getting out of hand. Faced with such problems, what are we do do?

Our first instinctive action is to look for someone with corresponding intellectual powers: we think of a Napoleon or an Archimedes. But detailed study of the distribution of man's intelligence shows that this method can give little. Figure 1, for instance, shows the distribution of the Intelligence Quotient in the

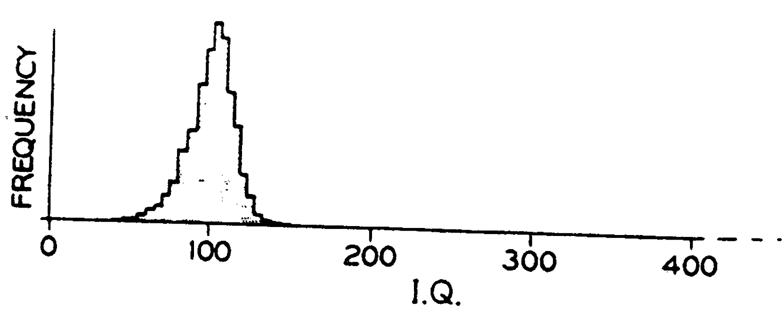

FIGURE 1. Distribution of the adult human Intelligence Quotient (after Wechsler, 1).

normal adult population, as found by Wechsler[1]. What is important for us now is not the shape on the left but the absolute emptiness on the right. A variety of tests by other workers have always yielded about the same result: a scarcity of people with I.Q.s over 150, and a total absence of I.Q.s over 200. Let us admit frankly that man's intellectual powers are as bounded as are those of his muscles. What then are we to do?

We can see something of how to proceed by comparing our position today in respect to intellectual problems with the position of the Romans in respect to physical problems. The Romans were doubtless often confronted by engineering and mechanical problems that demanded extreme physical strength. Doubtless the exceptionally strong slave was most useful, and doubtless the Romans sometimes considered the possibility of breeding slaves of even greater strength. Nevertheless, such plans were misdirected: only when men turned from their own powers to the powers latent in nature was the revolution inaugurated by Watt possible. Today, a workman comes to his task with a thousand horsepower available, though his own muscles will provide only about one-tenth. He gets this extra power by using a "power-amplifier". Had the present day brain-worker an "intelligence-amplifier" of the same ratio, he would be able to bring to his problems an I.Q. of a million.

If intellectual power is to be so developed, we must, somehow, construct amplifiers for intelligence – devices that, supplied with a little intelligence, will emit a lot. To see how this is to be done, let us look more closely at what is implied.

2. The Criterion of Intelligence

Let us first be clear about what we want. There is no intention here to inquire into the "real" nature of intelligence (whatever that may mean). The position is simple: we have problems and we want answers. We proceed then to ask, where are the answers to be found?

It has often been remarked that any random sequence, if long enough, will contain *all* the answers. Nothing prevents a child from doodling

$$\cos^2 x + \sin^2 x = 1,$$

or a dancing mote in the sunlight from emitting the same message in Morse or a similar code. Let us be more definite. If each of the above 13 symbols might have been any one of 50 letters and elementary signs, then as 50^{13} is approximately 2^{73}, the equation can be given in coded form by 73 binary symbols. Now consider a cubic centimeter of air as a turmoil of colliding molecules. A particular molecule's turnings after collision, sometimes to the left and sometimes to the right, will provide a series of binary symbols, each 73 of which, on some given code, either will or will not represent the equation. A simple calculation from the known facts shows that the molecules in every cubic centimeter of air are emitting this sequence *correctly* over a hundred thousand times a second. The objection that "such things don't happen" cannot stand.

Doodling, then, or any other random activity, is capable of producing all that is required. What spoils the child's claim to be a mathematician is that he will doodle, with equal readiness, such forms as

$$\cos^2 x + \sin^2 x = 2 \qquad \text{or} \qquad \text{ci)xsi} = \text{nx1}$$

or any other variation. After the child has had some mathematical experience he will stop producing these other variations. He becomes not more but less productive: he becomes selective.

The close, indeed essential, relation between intelligence and selection is shown clearly if we examine the tests specially devised for its objective measurement. Take, for instance, those of the Terman and Merrill[2] series for Year IV. In the first test the child is shown a picture of a common object and is asked to give its name. Out of all the words he knows he is asked to select one. In the second test, three model objects – motor-car, dog, show – are placed in a row and seen by the child; then all are hidden from him and a cover is placed over the dog; he is then shown motor-car, cover, shoe, and asked what is under the cover. Again his response is correct if, out of all possible words, he can select the appropriate one. Similarly the other tests, for all ages, evoke a response that is judged "correct" or "incorrect" simply by the subject's power of appropriate selection.

The same fact, that getting a solution implies selection, is shown with special clarity in the biological world. There the problems are all ultimately of how to achieve survival, and survival implies that the essential variables – the supply of food, water, etc. – are to be kept within physiological limits. The solutions to these problems are thus all selections from the totality of possibilities.

The same is true of the most important social and economic problems. What is wanted is often simple enough in aim – a way of ensuring food for all with an increasing population, or a way of keeping international frictions small in spite of provocations. In most of these problems the aim is the keeping of certain variables within assigned limits; and the problem is to find, amid the possibilities, some set of dynamic linkages that will keep the system both stable, and stable within those limits. Thus, finding the answer is again equivalent to achieving an appropriate selection.

The fact is that in admiring the *productivity* of genius our admiration has been misplaced. Nothing is easier than the generation of new ideas: with some suitable interpretation, a kaleidoscope, the entrails of a sheep, or a noisy vacuum tube will generate them in profusion. What is remarkable in the genius is the discrimination with which the possibilities are winnowed.

A possible method, then, is to use some random source for the generation of all the possibilities and to pass its output through some device that will select the answer. But before we proceed to make the device we must dispose of the critic who puts forward this well known argument: as the device will be made by some designer, it can select only what he has made it to select, so it can do no more than he can. Since this argument is clearly plausible, we must examine it with some care.

To see it in perspective, let us remember that the engineers of the middle ages, familiar with the principles of the lever and cog and pulley, must often have said that as no machine, worked by a man, could pull out more work than he put in, therefore no machine could ever amplify a man's power. Yet today we see one man keeping all the wheels in a factory turning by shovelling coal into a furnace. It is instructive to notice just how it is that today's stoker defeats the medieval engineer's dictum, while being still subject to the law of the conservation of energy. A little thought shows that the process occurs in two stages. In Stage One the stoker lifts the coal into the furnace; and over this stage energy is conserved strictly. The arrival of the coal in the furnace is then the beginning of Stage Two, in which again energy is conserved, as the burning of the coal leads to the generation of steam and ultimately to the turning of the factory's wheels. By making the whole process, from stoker's muscles to factory wheel, take place in two stages, involving two lots of energy whose sizes can vary with some independence, the modern engineer can obtain an overall amplification. Can we copy this method in principle so as to get an amplification in selection?

3. The Selection-amplifier

The essence of the stoker's method is that he uses his (small) power to bring into action that which will provide the main power. The designer, therefore, should use his (small) selectivity to bring into action that which is going to do the main selecting. Examples of this happening are common-place once one knows what to look for. Thus a garden sieve selects stones from soil; so if a gardener has sieves of different mesh, his act of selecting a sieve means that he is selecting, not the stones from the soil, but that which will do the selecting. The end result is that the stones are selected from the soil, and this has occurred as a consequence of his primary act; but he has achieved the selection mediately, in two stages. Again, when the directors of a large firm appoint a Manager of Personnel, who will attend to the selection of the staff generally, they are selecting that which will do the main selecting. When the whole process of selection is thus broken into two stages the details need only a little care for there to occur an amplification in the degree of selection exerted.

In this connection it must be appreciated that the *degree* of selection exerted is not defined by what is selected: it depends also on what the object is selected from. Thus, suppose I want to telephone for a plumber, and hesitate for a moment between calling Brown or Green, who are the only two I know. If I decide to ring up Green's number I have made a one-bit selection. My secretary, who will get the number for me, is waiting with directory in hand; she also will select Green's number, but she will select it from 50,000 other numbers, a 15.6-bit selection. (Since a 1-bit selection has directly determined a 15.6-bit selection, some aplification has occurred.) Thus two different selectors can select the same thing and yet exert quite different degrees of selection.

The same distinction will occur in the machine we are going to build. Thus, suppose we are tackling a difficult social and economic problem; we first select what we want, which might be:

An organization that will be stable at the conditions:

Unemployed	$<$	100,000 persons
Crimes of violence	$<$	10 per week
Minimal income per family	$>$	£500 per annum

This is *our* selection, and its degree depends of what other conditions we might have named but did not. The solving-machine now has to make *its* selection, finding this organization among the multitudinous other possibilities in the society. We and the solving-machine are selecting the same entity, but we are selecting it from quite different sets, or contexts, and the degrees of selection exerted can vary with some independence. (The similarity of this relation with those occurring in information theory is unmistakable; for in the latter the information-content of a message depends not only on what is in the message but on what population of message it came from[3,4].)

The building of a true selection-amplifier – one that selects over a greater range than that covered when it was designed – is thus possible. We can now proceed to build the system whose selectivity, and therefore whose intelligence, exceeds that of its designer.

(From now on we shall have to distinguish carefully between *two* problems: our problem, which is to design and build the solving-machine, and the solving-machine's problem – the one we want it to solve.)

4. Basic Design

Let us suppose for definiteness that the social and economic problem of the previous article is to be the solver's problem. How can we design a solver for it? The construction would in practice be a formidable task, but here we are concerned only with the principles. First, how is the selection to be achieved automatically?

SELECTION BY EQUILIBRIUM. We can take advantage of the fact that if any two determinate dynamic systems (X and S in Figure 2) are coupled through channels G and U so that each affects the other, then any resting state of the whole (that is, any state at which it can stay permanently), must be a resting state in each of the two parts individually, each being in the conditions provided by the other. To put it more picturesquely, each part has a power of veto over resting states proposed by the other. (The formulation can be made perfectly precise in the terms used in Article 6.)

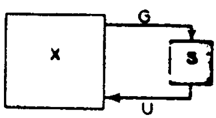

FIGURE 2

It is only a change of words to say that each part acts selectively towards the resting states of the other. So if S has been specially built to have resting states only on the occurrence of some condition ξ in S, then S's power of veto ensures that a resting state of the whole will always imply ξ in S. Suppose next that the linkage G is such that G will allow ξ to occur in S if and only if the condition η occurs in X. S's power of veto now ensures that any resting state of the whole must have condition η in X. So the selection of S and G to have these properties ensures that the only states in X that can be permanent are those that have the condition η.

It must be noticed that the selection of η, in the sense of its retention in X, has been done in two stages. The first occurred when the designer specified S and G and ξ. The second occurred when S, acting without further reference to the designer, rejected state after state of X, accepting finally one that gave the condition η in X. The designer has, in a sense, selected η, as an ultimate consequence of his actions, but his actions have worked through two stages, so the selectivity achieved in the second stage may be larger, perhaps much larger, than that used in the first.

The application of this method to the solving of the economic problem is, in principle, simple. We identify the real economic world with X and the conditions that we want to achieve in it with η. The selection of η in X is beyond our power, so we build, and couple to it, a system S, so built that it has a resting state if and only if its information through G is that η has occurred in a resting state in X. As time progresses, the limit of the whole system, X and S, is the permanent retention of η in X. The designer has to design and build S and G, and to couple it to X; after that the process occurs, so far as he is concerned, automatically.

5. The Homeostat

To see the process actually at work, we can turn briefly to the Hemeostat. Though it has been described fully elsewhere[5], a description of how its action appears in the terms used here may be helpful in illustration. (Figure 3 is intended to show its principle, not its actual appearance.)

It consists of four boxes (F) of components, freely supplied with energy, that act on one another in a complex pattern of feedbacks, providing initially a somewhat chaotic system, showing properties not unlike those sometimes seen in our own society. In this machine, S has been built, and G arranged, so that S has a resting state when and only when four needles N are stable at the central positions. These are the conditions η. N and the F's correspond to X in Figure 2. S affects the F's through the channel U, whose activity causes changes in the conditions within the boxes.

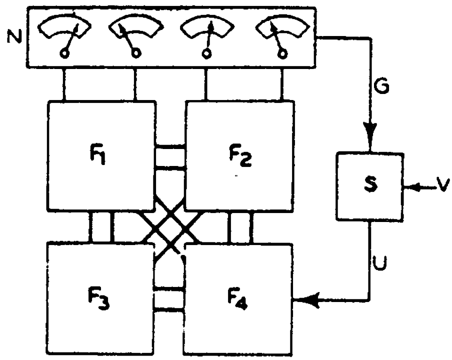

FIGURE 3

Suppose now that the conditions within the boxes, or in the connections between them, are set in some random way, say as a by-stander pleases; then if the conditions so set do not satisfy η, S goes into activity and enforces changes until η is restored. Since η may refer to certain properties of the stability within F, the system has been called "ultrastable", for it can regulate the conditions of its own stability within F. What is important in principle is that the combinations in F that restore η were not foreseen by the designer and programmer in detail; he provided only a random collection of about 300,000 combinations (from a table of random numbers), leaving it to S to make the detailed selection.

One possible objection on a matter of principle is that all the variation going to the trials in X seems, in Figure 2, to be coming from S and therefore from the designer, who has provided S. The objection can easily be met, however, and the alteration introduces an important technical improvement. To make the objection invalid, all the designer has to do is to couple S, as shown in Figure 3, to some convenient source of random variation V – a noisy vacuum tube say – so that the SV combination

(i) sends disturbance of inexhaustible variety along U if ξ is not occurring in S, and

(ii) keeps U constant, i.e., blocks the way from V to U, if ξ *is* occurring.

In the Hemeostat, V is represented by the table of random numbers which determined what entered F along U. In this way the whole system, X and S, has available the inexhaustible random variation that was suggested in Article 2 as a suitable source for the solutions.

6. Abstract Formulation

It is now instructive to view the whole process from another point of view, so as to bring out more clearly the deep analogy that exists between the amplification of power and that of intelligence.

Consider the engineer who has, say, some ore at the foot of a mine-shaft and who wants it brought to the surface. The power required is more than he can supply personally. What he does is to take some system that is going to change, by the laws of nature, from low entropy to high, and he couples this system to his ore, perhaps through pistons and ropes, so that "low entropy" is couple to "ore down" and "high entropy" to "ore up". He then lets the whole system go, confident that as the entropy goes from low to high so will it change the ore's position from down to up.

Abstractly (Figure 4) he has a process that is going, by the laws of nature, to pass from state H_1 to state H_2. He wants C_1 to change to C_2. So he couples H_1 to C_1 and H_2 to C_2. Then the system, in changing from H_1 to H_2, will change from C_1 to C_2, which is what he wants. The arrangement is clearly both necessary and sufficient.

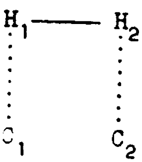

FIGURE 4

The method of getting the problem-solver to solve the set problem can now be seen to be of essentially the same form. The job to be done is the bring of X, in Figure 2, to a certain condition or "sollution" η. What the intelligence engineer does first to build a system, X and S, that has the tendency, by the laws of nature, to go to a state of equilibrium. He arranges the coupling between them so that "not at equilibrium" is coupled to not-η, and "at equilibrium" to η. He then lets the system go, confident that as the passage of time takes the whole to an equilibrium, so will the conditions in X have to change from not η to η. He does not make the conditions in X change by his own efforts, but allows the basic drive of nature to do the work.

This is the fundamental principle of our intelligence-amplifier. Its driving power is the tendency for entropy to increase, where "entropy" is used, not as understood in heat engines but as understood in stochastic processes.

AXIOMATIC STATEMENT. Since we are considering systems of extreme generality, the best representation of them is given in terms of the theory of sets. I use the concepts and terminology of Bourbaki[6].

From this point of view a machine, or any system that behaves in a determinate way, can be at any one of a set of states at a given moment. Let M be the set of states and μ some one of them. Time is assumed to be discrete, changing by unit intervals. The internal nature of the machine, whose details are irrelevant in this context, causes a transformation to occur in each interval of time, the state μ passing over determinately to some state μ' (not necessarily different from μ), thereby defining a mapping t of M in M:

$$t: \mu \to \mu' = t(\mu).$$

If the machine has an input, there will be a set I of input states ι, to each of which will correspond a mapping t_ι. The states ι may, of course, be those of some other machine, or may be determined by it; in this way machine may be coupled to machine. Thus if machine N with states ν has a set K of inputs κ and transformations u_κ, then machines M and N can be coupled by defining a mapping ζ of M in K, $\kappa = \zeta(\mu)$, and a mapping of m of N in I, $\iota = m(\nu)$, giving a system whose states are the couples (μ, ν) and whose changes with time are defined by the mapping, of M x N in M x N:

$$(\mu,\nu) \to ({}^{t}m(\nu)^{(\mu)}, {}^{u}\zeta(\mu)^{(\nu)}).$$

The abstract specification of the principle of ultrastability is as follows, Figure 5 corresponding to Figure 2:

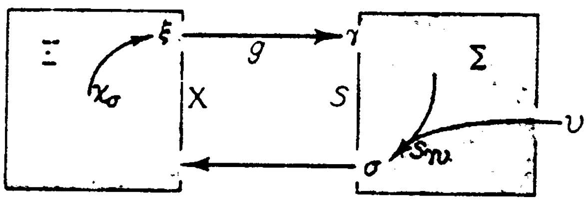

FIGURE 5

GIVEN:

(1) A set Γ consisting of two elements γ_1 and γ_2;
(2) A set Ξ of elements ξ;
(3) A mapping g of Ξ in Γ;

(4) A set Σ of elements σ;

(5) A family of mappings χ_σ of Ξ in Ξ;

(6) A random variable ν, with inexhaustible variety;

(7) A double family of mappings $s_{\gamma\nu}$ of Σ in Σ, with the property that, for all $\sigma \in \Sigma$ and all values of ν,

$$s_{\gamma_1}(\sigma) \neq \sigma \qquad \text{and} \qquad s_{\gamma_2}(\sigma) \neq \sigma;$$

(8) Time, advancing by discrete intervals, induces the operations χ_σ and $s_{\gamma\nu}$ and the successive values of ν simultaneously, once in each interval.

THEOREM: If the series of states of Ξ, induced by time, has a limit $\xi^* \in g^{-1}(\gamma_2)$.

PROOF. The state of the whole system, apart from ν, is given by the couple (ξ, σ), an element in $\Xi \times \Sigma$. The passage of one interval of time induces the mapping $(\xi, \sigma) \to \left(\chi_\sigma(\xi), s_{g(\xi),\nu}(\sigma) \right)$. If the series is at a limit-state (ξ^*, σ^*), then $s_{g(\xi^*),\nu}(\sigma^*) = \sigma^*$ for all values of ν. Therefore $g(\xi^*) = \gamma_2$, and $\xi^* \in g^{-1}(\gamma_2)$.

SECTION II

With this theorem our problem is solved, at least in principle. Systems so built can find solutions, and are not bounded by the powers of their designers. Nevertheless, there is still a long way to go before a man-made intelligence-amplifier will be actually in operation. Space prohibits any discussion of the many subsidiary problems that arise, but there is one objection that will probably be raised, suggesting that the method can never be of use, that can be dealt with here. It is of the greatest importance in the subject, so the remainder of the paper will be devoted to its consideration.

7. Duration of Trials

What we have discussed so far has related essentially to a process that, it is claimed, has "solution" as its limit. The question we are now coming to is, how fast is the convergence? how much time will elapse before the limit is reached?

A first estimate is readily made. Many of the problems require that the answer is an n-tuple, the solver being required to specify the value of each of n components. Thus the answer to an economic problem might require answeres to each of the questions:

(1) What is the optimal production-ratio of coal to oil?

(2) What amount should be invested annually in heavy industry?

(3) What should be the differential between the wages of skilled and unskilled workers?

except that in a real system n would be far larger than three.

A first estimate of the number of states to be searched can be made by finding the number of states in each component, assuming independence, and then finding the product so as to give the number of combinations. The time, with a state by state search, will then be proportional to this product. Thus, suppose that there are on a chessboard ten White and ten Black men; each can move to one of six squares; how many possibilities will have to be searched if I am to find the best next two moves, taking White's two moves and Black's two into account? With each man having about six moves (when captures are allowed for), the possible moves at each step are approximately 6^{10}; and the total possibilities are about 6^{40}. To find the best of them, even if some machine ran through them at a million a second, would take nearly a billion billion years – a prohibitively long time. The reason is that the time taken increases, in this estimate, exponentially with the number of components; and the exponential rate of increase is extremely fast. The calculation is not encouraging, but it is very crude; may it be seriously in error?

It is certainly in error to some extent, for it is not strictly an estimate but an upper bound. It will therefore always err by over-estimation. This however is not the chief reason for thinking that our method may yet prove practical. The reason for thinking this will be given in the next three articles.

8. The Method of Models

The first factor that can be used to reduce the time of search is of interest because it is almost synonymous with the method of science itself. It is, to conduct the search, not in the real physical thing itself but in a model of it, the model being chosen so that the search can proceed in it very much more rapidly. Thus Leverrier and Adams, searcing for a planet to explain the aberrations of Uranus, used pencil, paper and mathematics rather than the more obvious telescope; in that way they found Neptune in a few months where a telescopic search might have taken a lifetime.

The essence of the method is worth noticing explicitly. There is a set R, containing the solutions r, a subset of R; the task is to find a number of r. The method of models can be used if we can find some other set R' whose elements can be put into correspondence with those R in such a way that the elements (a set r') in R' that correspond to those in r can be recognized. The search is then conducted in R' for one of r'; when successful, the correspondence, used inversely, identifies a solution in R. For the method to be worth using, the search in R$'$ must be so much faster than that in R that the time taken in the three operations

(i) change from R to R$'$,
(ii) search in R$'$,
(iii) change back from r$'$ to r

is less than that taken in the single operation of searching in R.

Such models are common and are used widely. Pilot plants are used for trials rather than a complete workshop. Trials are conducted in the drawing-office rather than at the bench. The analogue computer is, of course, a model in this sense, and so, in a more subtle way, is the digital computer. Mathematics itself provides a vast range of models which can be handled on paper, or made to "behave", far faster than the systems to which they refer.

The use of models with the word extended to mean any structure isomorphic with that of the primary system, can thus often reduce the time taken to a fraction of what might at first seem necessary.

9. Constraints

A second reason why the time tends to fall below that of the exponential upper bound is that often the components are *not* independent, and the effect of this is always to reduce the range of possibilities. It will, therefore, other things being equal, reduce the time of search.

CONSTRAINT BY RELATION. Suppose we are looking for a solution in the n-tuple $(a_1, \ldots, a_n)$, where a_1 is an element in a set A_1, etc. The solution is then one of a subset of the product set $A_1 \times A_2 \times \ldots \times A_n$. A relation between the a's, $\phi(a_1, \ldots, a_n)$, always defines a subset of the product space [6], so if the relation holds over the a's, the solution will be found in the subset of the product-space defined by ϕ. An obvious example occurs when there exist invariants over the a's. k invariants can be used to eliminate k of the a's, which shows that the original range of variation was over n − k, not over n, dimensions. More generally, every law, whether parliamentary, or natural, or algebraic[7] is a constraint, and acts to narrow the range of variation and, with it, the time of search.

Similarly, every "entity" that can be recognized in the range of variation holds its individuality only if its parts do *not* vary over the full range conceivable. Thus, a "chair" is recognizable as a thing partly because its four legs do not move in space with all the degrees of freedom possible to four independent objects. The fact that their actual degrees of freedom are 6 instead of 24 is a measure of their cohesion. Conversely, their cohesion implies that any reckoning of possibilities must count 6 dimensions in the product or phase space, not 24.

It will be seen therefore that ***every relation that holds between the components of an n-tuple lessens the region of search.***

CONSTRAINT BY CONTINUITY. Another common constraint on the possibilities occurs where there are functional relations within the system such that the function is continuous. Continuity is a restriction, for if $y = f(z)$ and f is continuous and a series of arguments z, z' and z'', . . . has the limit z^*, then the corresponding series y, y', y'' . . . must have the limit $f(z^*)$. Thus f is not free to relate y and z, y' and z', y'' and z'', . . . arbitrarily, as it could do if it were unrestricted. This fact can also be expressed by saying that as adjacent values of z make the values of $f(z)$ adjacent, these values of $f(z)$ tend to be highly correlated, so that the values of $f(z)$ can be adequately explored by a mere sampling of the possibilities in z: the values of y do not have to be tested individually.

A rigorous discussion of the subject would lead into the technicalities of topology; here it is sufficient to notice that the continuity of f puts restrictions on the range of possibilities that will have to be searched for a solution.

Continuity helps particularly to make the search easy when the problem is to find what values of a, b, c, . . . will make some function $\lambda(a,b,c, \ldots)$ a maximum. Wherever λ is discontinuous and arbitrary there is no way of finding the maximum, or optimum, except by trying every combination of the arguments individually; but where λ is continuous a maximum can often be proceded to directly. The thesis is well illustrated in the art of aircraft design, which involves knowing the conditions that will make the strength, lightness, etc., an optimum. In those aspects of design in which the behavior of the aircraft is a continuous function of the variables of design, the finding of an optimum is comparatively direct and rapid; where however the relations are discontinuous, as happens at the critical values of the Reynolds' and Mach numbers, then the finding of an optimum is more laborious.

Chess, in fact, is a difficult game largely because of the degree to which it is discontinuous, in the sense that the "value" of a position, to White say, is by no means a continuous function of the positions of the component pieces. If a rook, for instance, is moved square by square up a column, the successive values of the positions vary, sometimes more or less continuously, but often with marked discontinuity. The high degree of discontinuity occurring throughout a game of chess makes it very unlike the more "natural" systems, in which continuity is common. With this goes the corollary that the test so often proposed for a mechanical brain – that it should play chess – may be misleading, in that it is by no means representative of the class of problem that a real mechanical brain will one day have to deal with.

The commonness of continuity in the world around us has undoubtedly played an important part in Darwinian evolution in that progress in evolution would have been far slower had not continuity been common. In this paper I have tended to stress the analogy of the solving process with that of the amplification of physical power I could equally have stressed its deep analogy with the processes of evolution, for there is the closest formal similarity between the process by which adaptation is produced automatically by Darwinian selection and the process by which a solution is produced automatically by mechanical selection of the type considered in Article 4. Be that as it may, every stockbreeder knows that selection for good strains can procede much more rapidly when the relation between genotype and phenotype is continuous[8]. The time taken for a given degree of improvement to be achieved is, of course, correspondingly reduced.

To sum up: Continuity being common in the natural world, the time taken in the solution of problems coming from it may be substantially below that given by the exponential bound.

CONSTRAINT BY PRIOR KNOWLEDGE. It is also only realistic to consider, in this connection, the effect on solving of knowledge accumulated in the past. Few problems are wholly new, and it is merely common sense that we, and the problem-solver, should make use of whatever knowledge has already been won.

The effect of such knowledge is again to put a constraint on the possibilities, lessening the regions that have to be searched, for past experience will act essentially by warning us that a solution is unlikely to lie in certain regions. The constraint is most marked when the problem is one of a class, of which several have already been solved. Having such knowledge about other members of the same class is equivalent to starting the solving process at some point that is already partly advanced towards the goal. Such knowledge can naturally be used to shorten the search.

"Solving a problem" can in fact be given a perfectly general representation. The phase- or sample-space of possibilities contains many points, most of them corresponding to "no solution" but a few of them corresponding to an acceptable "solution". Finding a solution then becomes a matter of starting somewhere in this unknown distribution of states and trying to find one of the acceptable states.

If the acceptable states are distributed wholly at random, i.e., with no recognizable pattern, then we are considering the case of the problem about which nothing, absolutely nothing, is known. After 200 years of scientific activity, such problems today are rare; usually some knowledge is available from past experience, and this knowledge can be used to constrain the region of search, making it smaller and the search more rapid. Thus suppose, as a simple example, that there are 114 states to be examined and that 40 of them are acceptable, i.e., correspond to solutions. If they are really scattered at random, as are the black squares in I of Figure 6, then the searcher has no better resource than to start somewhere and to wander at random. With this method, in this particular example he will usually require about 3.6 trials to find a black square. If, however, past experience has shown that the black squares are distributed as in II, advantage can be taken to shorten this search; for wherever the search may start, the occurrence of two white squares in succession shows the searcher that he is in an all-white quadrant. A move of four squares will take him to a checkered quadrant where he will find a dark square either at once or on the next move. In this case his average number of trials need not exceed about 2.4 (if we ignore the complications at the boundaries).

I FIGURE 6 II

Information theory and the strategy of games are both clearly involved here, and an advanced knowledge of such theories will undoubtedly be part of the technical equipment of those who will have to handle the problem-solvers of the future. Meanwhile we can sum up this Article by saying that it has shown several factors that tend to make the time of search less than that given by the exponential bound.

10. Selection by Components

We now come to what is perhaps the most potent factor of all for reducing the time taken. Let us assume, as before, that the solver is searching for an element in a set and that the elements are n-tuples. The solver searches, then, for the n values of the n components.

In the most "awkward" problems there is no means of selection available other than the testing of every element. Thus, if each of the n components can take any one of k values, the number of states to be searched is k^n, and the time taken will be proportional, as we saw in Article 7, to the exponential bound. This case, however, is the worst possible.

Next let us consider the other extreme. In this case the selection can be conducted component by component, each component being identified independently of the others. When this occurs, a very great reduction occurs in the time taken, for now the number of states to be searched is kn.

This second number may be very much smaller than the first. To get some idea of the degree of change implied, let us take a simple example. Suppose there are a thousand components, each of which has a hundred possibilities. If the search has to be over every *combination* separately, the number of states to be searched is the exponential bound, 100^{1000}. If the states could be examined at one a second it would take about 10^{1993} years, a duration almost beyond thinking. If we tried to get through the same selection by the method of models, using some device that could test, say, a million in each microsecond, the time would drop to 10^{1981} years, which is practically no change at all. If however the components could be selected individually and independently, then the number of selections would drop to 100,000; and at one per second the problem could be solved in a little over a day.

This illustration will suffice to illustrate the general rule that selection by components is so much faster than selection by elements that a search that is utterly impractical by the one may be quite easy by the other.

The next questions is to what extent problems are susceptible of this faster method. About *all* problems one can say nothing, for the class is not defined, but about the problems of social and economic type, to which this paper is specially directed, it is possible to form some estimate of what is to be expected. The social and economic system is highly dynamic, active by its own energies, and controllable only through a number of parameters. The solver thus looks for suitable values in the n-tuple of parameters. (Its number of components will not usually be the same as those of the variables mentioned in the next paragraph.)

Now a dynamic system may be "reducible"; this happens when the system that seems to be single really consists of two or more adjacent, but functionally independent, parts. If the parts, P and Q say, are quite independent, the whole is "completely" reducible; if part P affects part Q, but Q has no effect on P, then the whole is simple "reducible". In the equations of such a system, which can be given in the canonical form[5]

$$\frac{dx_1}{dt} \doteq f_1(x_1, \ldots, x_n)$$

$$\frac{dx_n}{} = f_n(x_1, \ldots, x_n)$$

reducibility is shown by the Jacobian of the n functions f_1 with respect to the n variables x_j having, when partitioned into the variables that belong to P and those that belong to Q, one or more all-zero quadrants, being like I when reducible and like II when completely reducible:

$$\underset{\text{I}}{\begin{bmatrix} J_1 & O \\ K & J_2 \end{bmatrix}} \qquad\qquad \underset{\text{II}}{\begin{bmatrix} J_1 & O \\ O & J_2 \end{bmatrix}}$$

Such reducibility corresponds, in a real dynamic system, with the possibility of searching to some extent by components, the components being the projections[6] of the abstract space on those variables, or "coordinates", that occur together in one part (P, or Q) of the reducible system. The more a system is reducible, the more does it offer the possibility of the search being made by the quick method of finding the components separately.

Though the question has not yet been adequately explored, there is reason to believe that part-functions, i.e., variables whose f_1 (above) are zero for many values of their arguments, are common in social and economic systems. They are ubiquitous in the physical world, as I have described elsewhere[5]. Whenever they occur, they introduce zeros into the Jacobian of the system (for if x_j is constant, over some interval of time, f_1 must be zero and therefore so will be $\delta f_1 / \delta x_j$); they therefore tend to introduce temporary reducibilities. This means that the problem may be broken up into a series of conditional and transient sub-problems, each of which has a solution that can be found more or less easily. The whole search thus has something of the method of search by components.

In this connection it is instructive to consider an observation of Shannon's[9] on the practicability of using relays as devices for switching, for it involves a property closely related to that of reducibility. One of the problems he solved was to find how many elements each relay would have to operate if the network was to be capable of realizing all possible functions on n variables. The calculation gave a number that was, apparently, ridiculously high for, did one not know, it would suggest that relays were unsuitable for practical use in switching. The apparent discrepancy proved to be due to the fact that the functions commonly required in switching are not as complicated as the class that can be considered in theory. The more they *look* complicated the more they tend to have hidden simplicities. These simplicities are of the form in which the function is "separable", , that is to say, the variables go in sets, much as the variables in a reducible system go in sets. Separability thus makes what seems to be an impractical number of elements become practical. It is not unlikely that reducibility will act similarly towards the time of search.

11. The Lower Bound

It can now be seen that problem-solving, as a process of selection, is related to the process of message-receiving as treated in information theory[3, 10]. The connection can be seen most readily by imagining that the selection of one object from N is to be made by an agent A who acts according to instructions received from B. As the successive elements appear, B will issue signals: ". . . , reject, reject, . . . , accept" thereby giving information in calculable quantity to A. Usually the number of binary signals so given is likely to be greater than the number necessary, for the probabilities of the two signals, "reject" and "accept", are by no means equal. The most efficient method of selection is that which makes the two signals equally likely. This will happen in the whole set of N can be dichotomized at each set of selection. In this way we can find a *lower bound* to the time taken by the process, which will be proportional to log N. This time is none other than the least time in which, using binary notation, the solution can be written down, that is, identified from its alternatives.

We see therefore that, though the upper (exponential) bould is forbiddingly high, the lower bound is reassuringly low. What time will actually be taken in some particular problem can only be estimated after a direct study of the particular problem and of the resources available. I hope, however, that I have said enough to show that the mere mention of the exponential bound is not enough to discredit the method proposed here. The possibility that the method will work is still open.

Summary

The question is considered whether it is possible for human constructors to build a machine that can solve problem: of more than human difficulty. If physical power can be amplified, why not intellectual?

Considerations show that:

Getting an answer to a problem is essentially a matter of selection.

Selection can be amplified.

A system with a selection-amplifier can be more selective than the man who built it.

Such a system is, in principle, capable of solving problems, perhaps in the social and economic world, beyond the intellectual powers of its designer.

A first estimate of the time it will take to solve a difficult problem suggests that the time will be excessively long: closer examination shows that the estimate is biased, being an upper bound.

The lower bound, achievable in some cases, is the time necessary for the answer to be written down in binary notation.

It is not impossible that the method may be successful in those social and economic problems to which the paper is specially addressed.

Bibliography

1 WECHSLER, D., *Management of Adult Intelligence.* Baltimore, Williams & Wilkins, 3rd Ed., 1944.

2 TERMAN, L.M., and MERRILL, M.A., *Measuring Intelligence.* London, Harrap & Co., 1937.

3 SHANNON, C.E. and WEAVER, W., *The Mathematical Theory of Communication.* Urbana, University of Illinois Press, 1949.

4 ASHBY, W. ROSS, *Can A Mechanical Chess-player Outplay its Designer?* Brit. J. Phil. Sci., 3, 44: 1952.

5. ASHBY, W. ROSS, *Design for a Brain.* London, Chapman & Hall; New York, John Wiley & Sons, 1952.

6 BOURBAKI, N., "Theorie des Ensembles." *A.S.E.I.* 1141. Paris, Herman et Cie, 2nd ed., 1951.

BOURBAKI, N. "Structures Algebriques." *A.S.E.I.* 1144. Paris, Hermann et Cie, 2nd ed., 1951.

8 LERNER, I.M. *Population Genetics and Animal Improvement.* Cambridge University Press, 1950.

9 SHANNON, C.E., "Synthesis of Two-terminal Switching Circuits." *Bell System tech. J.,* 28, 59-98, 1949.

10 WIENER, N., *Cybernetics.* New York, John Wiley & Sons, 1948.