IX On Learning and Self-Reproducing Machines

Two of the phenomena which we consider to be characteristic of living systems are the power to learn and the power to reproduce themselves. These properties, different as they appear, are intimately related to one another. An animal that learns is one which is capable of being transformed by its past environment into a different being and is therefore adjustable to its environment within its individual lifetime. An animal that multiplies is able to create other animals in its own likeness at least approximately, although not so completely in its own likeness that they cannot vary in the course of time. If this variation is itself inheritable, we have the raw material on which natural selection can work. If the hereditary invariability concerns manners of behavior, then among the varied patterns of behavior which are propagated some will be found advantageous to the continuing existence of the race and will establish themselves, while others which are detrimental to this continuing existence will be eliminated. The result is a certain sort of racial or phylogenetic learning, as contrasted with the ontogenetic learning of the individual. Both ontogenetic and phylogenetic learning are modes by which the animal can adjust itself to its environment.

Both ontogenetic and phylogenetic learning, and certainly the latter, extend themselves not only to all animals but to plants and, indeed, to all organisms which in any sense may be considered to be living. However, the degree to which these two forms of learning are found to be important in different sorts of living beings varies widely. In man, and to a lesser extent in the other mammals, ontogenetic learning and individual adaptability are raised to the highest point. Indeed, it may be said that a large part of the phylogenetic learning of man has been devoted to establishing the possibility of good ontogenetic learning.

It has been pointed out by Julian Huxley in his fundamental paper on the mind of birds[1] that birds have a small capacity for ontogenetic learning. Something similar is true in the case of insects, and in both instances it may be associated with the terrific demands made on the individual by flight and the consequential pre-emption of the capabilities of the nervous system which might otherwise be applied to ontogenetic learning. Complicated as the behavior patterns of birds are—in flying, in courtship, in the care of the young, and in nest building—they are carried out correctly the very first time without the need of any large amount of instruction from the mother.

It is altogether appropriate to devote a chapter of this book to these two related subjects. Can man-made machines learn and can they reproduce themselves? We shall try to show in this chapter that in fact they can learn and can reproduce themselves, and we shall give an account of the technique needed for both these activities.

The simpler of these two processes is that of learning, and it is there that the technical development has gone furthest. I shall talk here particularly of the learning of game-playing machines which enables them to improve the strategy and tactics of their performance by experience.

There is an established theory of the playing of games—the von Neumann theory.[2] It concerns a policy which is best considered by working from the end of the game rather than from the beginning. In the last move of the game, a player strives to make a winning move if possible, and if not, then at least a drawing move. His opponent, at the previous stage, strives to make a move which will prevent the other player from making a winning or a drawing move. If he can himself make a winning move at that stage, he will do so, and this will not be the next to the last but the last stage of the game. The other player at the move before this will try to act in such a way that the very best resources of his opponent will not prevent him from ending with a winning move, and so on backward.

There are games such as ticktacktoe where the entire strategy is known, and it is possible to start this policy from the very beginning. When this is feasible, it is manifestly the best way of playing the game. However, in many games like chess and checkers our knowledge is not sufficient to permit a complete strategy of this sort, and then we can only approximate to it. The von Neumann type of approximate theory tends to lead a player to act with the utmost caution, assuming that his opponent is the perfectly wise sort of a master.

This attitude, however, is not always justified. In war, which is a sort of game, this will in general lead to an indecisive action which will often be not much better than a defeat. Let me give two historical examples. When Napoleon fought the Austrians in Italy, it was part of his effectiveness that he knew the Austrian mode of military thought to be hidebound and traditional, so that he was quite justified in assuming that they were incapable of taking advantage of the new decision-compelling methods of war which had been developed by the soldiers of the French Revolution. When Nelson fought the combined fleets of continental Europe, he had the advantage of fighting with a naval machine which had kept the seas for years and which had developed methods of thought of which, as he was well aware, his enemies were incapable. If he had not made the fullest possible use of this advantage, instead of acting as cautiously as he would have had to act under the supposition that he was facing an enemy of equal naval experience, he might have won in the long run but could not have won so quickly and decisively as to establish the tight naval blockade which was the ultimate downfall of Napoleon. Thus, in both cases, the guiding factor was the known record of the commander and of his opponents, as exhibited statistically in the past of their actions, rather than an attempt to play the perfect game against the perfect opponent. Any direct use of the von Neumann method of game theory in these cases would have proved futile.

In a similar way, books on chess theory are not written from the von Neumann point of view. They are compendia of principles drawn from the practical experience of chess players playing against other chess players of high quality and wide knowledge; and they establish certain values or weightings to be given to the loss of each piece, to mobility, to command, to development, and to other factors which may vary with the stage of the game.

It is not very difficult to make machines which will play chess of a sort. The mere obedience to the laws of the game, so that only legal moves are made, is easily within the power of quite simple computing machines. Indeed, it is not hard to adapt an ordinary digital machine to these purposes.

Now comes the question of policy within the rules of the game. Every evaluation of pieces, command, mobility, and so forth, is intrinsically capable of being reduced to numerical terms; and when this is done, the maxims of a chess book may be used for the determination of the best moves of each stage. Such machines have been made; and they will play a very fair amateur chess, although at present not a game of master caliber.

Imagine yourself in the position of playing chess against such a machine. To make the situation fair, let us suppose you are playing correspondence chess without the knowledge that it is such a machine you are playing and without the prejudices that this knowledge may excite. Naturally, as always is the case with chess, you will come to a judgment of your opponent's chess personality. You will find that when the same situation comes up twice on the chessboard, your opponent's reaction will be the same each time, and you will find that he has a very rigid personality. If any trick of yours will work, then it will always work under the same conditions. It is thus not too hard for an expert to get a line on his machine opponent and to defeat him every time.

However, there are machines that cannot be defeated so trivially. Let us suppose that every few games the machine takes time off and uses its facilities for another purpose. This time, it does not play against an opponent, but examines all the previous games which it has recorded on its memory to determine what weighting of the different evaluations of the worth of pieces, command, mobility, and the like, will conduce most to winning. In this way, it learns not only from its own failures but its opponent's successes. It now replaces its earlier valuations by the new ones and goes on playing as a new and better machine. Such a machine would no longer have as rigid a personality, and the tricks which were once successful against it will ultimately fail. More than that, it may absorb in the course of time something of the policy of its opponents.

All this is very difficult to do in chess, and as a matter of fact the full development of this technique, so as to give rise to a machine that can play master chess, has not been accomplished. Checkers offers an easier problem. The homogeneity of the values of the pieces greatly reduces the number of combinations to be considered. Moreover, partly as a consequence of this homogeneity, the checker game is much less divided into distinct stages than the chess game. Even in checkers, the main problem of the end game is no longer to take pieces but to establish contact with the enemy so that one is in a position to take pieces. Similarly, the valuation of moves in the chess game must be made independently for the different stages. Not only is the end game different from the middle game in the considerations which are paramount, but the openings are much more devoted to getting the pieces into a position of free mobility for attack and defense than is the middle game. The result is that we cannot be even approximately content with a uniform evaluation of the various weighting factors for the game as a whole, but must divide the learning process into a number of separate stages. Only then can we hope to construct a learning machine which can play master chess.

The idea of a first-order programming, which may be linear in certain cases, combined with a second-order programming, which uses a much more extensive segment of the past for the determination of the policy to be carried out in the first-order programming, has been mentioned earlier in this book in connection with the problem of prediction. The predictor uses the immediate past of the flight of the airplane as a tool for the prediction of the future by means of a linear operation; but the determination of the correct linear operation is a statistical problem in which the long past of the flight and the past of many similar flights are used to give the basis of the statistics.

The statistical studies necessary to use a long past for a determination of the policy to be adopted in view of the short past are highly non-linear. As a matter of fact, in the use of the Wiener-Hopf equation for prediction,[3] the determination of the coefficients of this equation is carried out in a non-linear manner. In general, a learning machine operates by non-linear feedback. The checker-playing machine described by Samuel[4] and Watanabe[5] can learn to defeat the man that programmed it in a fairly consistent way on the basis of from 10 to 20 operating hours of programming.

Watanabe's philosophical ideas on the use of programming machines are very exciting. On the one hand, he is treating a method of proving an elementary geometrical theorem which shall conform in an optimal way according to certain criteria of elegance and simplicity, as a learning game to be played not against an individual opponent but against what we may call "Colonel Bogey." A similar game which Watanabe is studying is played in logical induction, when we wish to build up a theory which is optimal in a similar quasi-aesthetic way, on the basis of an evaluation of economy, directness, and the like, by the determination of the evaluation of a finite number of parameters which are left free. This, it is true, is only a limited logical induction, but it is well worth studying.

Many forms of the activity of struggle, which we do not ordinarily consider as games, have a great deal of light thrown on them by the theory of game-playing machines. One interesting example is the fight between a mongoose and a snake. As Kipling points out in "Rikki-Tikki-Tavi," the mongoose is not immune to the poison of the cobra, although it is to some extent protected by its coat of stiff hairs which makes it difficult for the snake to bite home. As Kipling states, the fight is a dance with death, a struggle of muscular skill and agility. There is no reason to suppose that the individual motions of the mongoose are faster or more accurate than those of the cobra. Yet the mongoose almost invariably kills the cobra and comes out of the contest unscathed. How is it able to do this?

I am here giving an account which appears valid to me, from having seen such a fight, as well as motion pictures of other such fights. I do not guarantee the correctness of my observations as interpretations. The mongoose begins with a feint, which provokes the snake to strike. The mongoose dodges and makes another such feint, so that we have a rhythmical pattern of activity on the part of the two animals. However, this dance is not static but develops progressively. As it goes on, the feints of the mongoose come earlier and earlier in phase with respect to the darts of the cobra, until finally the mongoose attacks when the cobra is extended and not in a position to move rapidly. This time the mongoose's attack is not a feint but a deadly accurate bite through the cobra's brain.

In other words, the snake's pattern of action is confined to single darts, each one for itself, while the pattern of the mongoose's action involves an appreciable, if not very long, segment of the whole past of the fight. To this extent the mongoose acts like a learning machine, and the real deadliness of its attack is dependent on a much more highly organized nervous system.

As a Walt Disney movie of several years ago showed, something very similar happens when one of our western birds, the road runner, attacks a rattlesnake. While the bird fights with beak and claws, and a mongoose with its teeth, the pattern of activity is very similar. A bullfight is a very fine example of the same thing. For it must be remembered that the bullfight is not a sport but a dance with death, to exhibit the beauty and the interlaced coordinating actions of the bull and the man. Fairness to the bull has no part in it, and we can leave out from our point of view the preliminary goading and weakening of the bull, which have the purpose of bringing the contest to a level where the interaction of the patterns of the two participants is most highly developed. The skilled bullfighter has a large repertory of possible actions, such as the flaunting of the cape, various dodges and pirouettes, and the like, which are intended to bring the bull into a position in which it has completed its rush and is extended at the precise moment that the bullfighter is ready to plunge the *estoque* into the bull's heart.

What I have said concerning the fight between the mongoose and the cobra, or the toreador and the bull, will also apply to physical contests between man and man. Consider a duel with the small-sword. It consists of a sequence of feints, parries, and thrusts, with the intention on the part of each participant to bring his opponent's sword out of line to such an extent that he can thrust home without laying himself open to a double encounter. Again, in a championship game of tennis, it is not enough to serve or return the ball perfectly as far as each individual stroke is considered; the strategy is rather to force the opponent into a series of returns which put him progressively in a worse position until there is no way in which he can return the ball safely.

These physical contests and the sort of games which we have supposed the game-playing machine to play both have the same element of learning in terms of experience of the opponent's habits as well as one's own. What is true of games of physical encounter is also true of contests in which the intellectual element is stronger, such as war and the games which simulate war, by which our staff officers win the elements of their military experience. This is true for classical war both on land and at sea, and is equally true with the new and as yet untried war with atomic weapons. Some degree of mechanization, parallel to the mechanization of checkers by learning machines, is possible in all these.

There is nothing more dangerous to contemplate than World War III. It is worth considering whether part of the danger may not be intrinsic in the unguarded use of learning machines. Again and again I have heard the statement that learning machines cannot subject us to any new dangers, because we can turn them off when we feel like it. But can we? To turn a machine off effectively, we must be in possession of information as to whether the danger point has come. The mere fact that we have made the machine does not guarantee that we shall have the proper information to do this. This is already implicit in the statement that the checker-playing machine can defeat the man who has programmed it, and this after a very limited time of working in. Moreover, the very speed of operation of modern digital machines stands in the way of our ability to perceive and think through the indications of danger.

The idea of non-human devices of great power and great ability to carry through a policy, and of their dangers, is nothing new. All that is new is that now we possess effective devices of this kind. In the past, similar possibilities were postulated for the techniques of magic, which forms the theme for so many legends and folk tales. These tales have thoroughly explored the moral situation of the magician. I have already discussed some aspects of the legendary ethics of magic in an earlier book entitled *The Human Use of Human Beings*.[6] I here repeat some of the material which I have discussed there, in order to bring it out more precisely in its new context of learning machines.

One of the best-known tales of magic is Goethe's "The Sorcerer's Apprentice." In this, the sorcerer leaves his apprentice and factotum alone with the chore of fetching water. As the boy is lazy and ingenious, he passes the work over to a broom, to which he has uttered the words of magic which he has heard from his master. The broom obligingly does the work for him and will not stop. The boy is on the verge of being drowned out. He finds that he has not learned, or has forgotten, the second incantation which is to stop the broom. In desperation, he takes the broomstick, breaks it over his knee, and finds to his consternation that each half of the broom continues to fetch water. Luckily, before he is completely destroyed, the master returns, says the Words of Power to stop the broom, and administers a good scolding to the apprentice.

Another story is the Arabian Nights tale of the fisherman and the genie. The fisherman has dredged up in his net a jug closed with the seal of Solomon. It is one of the vessels in which Solomon has imprisoned the rebellious genie. The genie emerges in a cloud of smoke, and the gigantic figure tells the fisherman that, whereas in his first years of imprisonment he had resolved to reward his rescuer with power and fortune, he has now decided to slay him out of hand. Luckily for himself, the fisherman finds a way to talk the genie back into the bottle, upon which he casts the jar to the bottom of the ocean.

More terrible than either of these two tales is the fable of the monkey's paw, written by W. W. Jacobs, an English writer of the beginning of the century. A retired English workingman is sitting at his table with his wife and a friend, a returned British sergeant-major from India. The sergeant-major shows his hosts an amulet in the form of a dried, wizened monkey's paw. This has been endowed by an Indian holy man, who has wished to show the folly of defying fate, with the power of granting three wishes to each of three people. The soldier says that he knows nothing of the first two wishes of the first owner, but the last one was for death. He himself, as he tells his friends, was the second owner but will not talk of the horror of his own experiences. He casts the paw into the fire, but his friend retrieves it and wishes to test its powers. His first is for £200. Shortly thereafter there is a knock at the door, and an official of the company by which his son is employed enters the room. The father learns that his son has been killed in the machinery, but that the company, without recognizing any responsibility or legal obligation, wishes to pay the father the sum of £200 as a solatium. The grief-stricken father makes his second wish—that his son may return—and when there is another knock at the door and it is opened, something appears which, we are not told in so many words, is the ghost of the son. The final wish is that this ghost should go away.

In all these stories the point is that the agencies of magic are literal-minded; and that if we ask for a boon from them, we must ask for what we really want and not for what we think we want. The new and real agencies of the learning machine are also literal-minded. If we program a machine for winning a war, we must think well what we mean by winning. A learning machine must be programmed by experience. The only experience of a nuclear war which is not immediately catastrophic is the experience of a war game. If we are to use this experience as a guide for our procedure in a real emergency, the values of winning which we have employed in the programming games must be the same values which we hold at heart in the actual outcome of a war. We can fail in this only at our immediate, utter, and irretrievable peril. We cannot expect the machine to follow us in those prejudices and emotional compromises by which we enable ourselves to call destruction by the name of victory. If we ask for victory and do not know what we mean by it, we shall find the ghost knocking at our door.

So much for learning machines. Now let me say a word or two about self-propagating machines. Here both the words *machine* and *self-propagating* are important. The machine is not only a form of matter, but an agency for accomplishing certain definite purposes. And self-propagation is not merely the creation of a tangible replica; it is the creation of a replica capable of the same functions.

Here, two different points of view come into evidence. One of these is purely combinatorial and concerns the question whether a machine can have enough parts and sufficiently complicated structure to enable self-reproduction to be among its functions. This question has been answered in the affirmative by the late John von Neumann. The other question concerns an actual operative procedure for building self-reproducing machines. Here I shall confine my attentions to a class of machines which, while it does not embrace all machines, is of great generality. I refer to the non-linear transducer.

Such machines are apparatuses which have as an input a single function of time and which have as their output another function of time. The output is completely determined by the past of the input; but in general, the adding of inputs does not add the corresponding outputs. Such pieces of apparatus are known as transducers. One property of all transducers, linear or non-linear, is an invariance with respect to a translation in time. If a machine performs a certain function, then, if the input is shifted back in time, the output is shifted back by the same amount.

Basic to our theory of self-reproducing machines is a canonical form of the representation of non-linear transducers. Here the notions of impedance and admittance, which are so essential in the theory of linear apparatus, are not fully appropriate. We shall have to refer to certain newer methods of carrying out this representation, methods developed partly by me[7] and partly by Professor Dennis Gabor[8] of the University of London.

While both Professor Gabor's methods and my own lead to the construction of non-linear transducers, they are linear to the extent that the non-linear transducer is represented with an output which is the sum of the outputs of a set of non-linear transducers with the same input. These outputs are combined with varying linear coefficients. This allows us to employ the theory of linear developments in the design and specification of the non-linear transducer. And in particular, this method allows us to obtain coefficients of the constituent elements by a least-square process. If we join this to a method of statistically averaging over the set of all inputs to our apparatus, we have essentially a branch of the theory of orthogonal development. Such a statistical basis of the theory of non-linear transducers can be obtained from an actual study of the past statistics of the inputs used in each particular case.

This is a rough account of Professor Gabor's methods. While mine are essentially similar, the statistical basis for my work is slightly different.

It is well known that electrical currents are not conducted continuously but by a stream of electrons which must have statistical variations from uniformity. These statistical fluctuations can be represented fairly by the theory of the Brownian motion, or by the similar theory of shot effect or tube noise, about which I am going to say something in the next chapter. At any rate, apparatus can be made to generate a standardized shot effect with highly specific statistical distribution, and such apparatus is being manufactured commercially. Note that tube noise is in a sense a universal input in that its fluctuations over a sufficiently long time will sooner or later approximate to any given curve. This tube noise possesses a very simple theory of integration and averaging.

In terms of the statistics of tube noise, we can easily determine a closed set of normal and orthogonal non-linear operations. If the inputs subject to these operations have the statistical distribution appropriate to tube noise, the average product of the output of two component pieces of our apparatus, where this average is taken with respect to the statistical distribution of tube noise, will be zero. Moreover, the mean square output of each apparatus can be normalized to one. The result is that the development of the general non-linear transducer in terms of these components results from an application of the familiar theory of orthonormal functions.

To be specific, our individual pieces of apparatus give outputs which are products of Hermite polynomials in the Laguerre coefficients of the past of the input. This is presented in detail in my *Nonlinear Problems in Random Theory*.

It is of course difficult to average in the first instance over a set of possible inputs. What makes this difficult task realizable is that the shot-effect inputs possess the property known as metric transitivity, or the ergodic property. Any integrable function of the parameter of distribution of shot-effect inputs has in almost every instance a time average equal to its average over the ensemble. This permits us to take two pieces of apparatus with a common shot-effect input, and to determine the average of their product over the entire ensemble of the possible inputs, by taking their product and averaging it over the time. The repertory of operations needed for all these processes involves nothing more than the addition of potentials, the multiplication of potentials, and the operation of averaging over time. Devices exist for all these. As a matter of fact, the elementary devices needed in Professor Gabor's methodology are the same as those needed in mine. One of his students has invented a particularly effective and inexpensive multiplying device depending on the piezoelectric effect on a crystal of the attraction of two magnetic coils.

What this amounts to is that we can imitate any unknown non-linear transducer by a sum of linear terms, each of fixed characteristics and with an adjustable coefficient. This coefficient can be determined as the average product of the outputs of the unknown transducer and a particular known transducer, when the same shot-effect generator is connected to the input of both. What is more, instead of computing this result on the scale of an instrument and then transferring it by hand to the appropriate transducer, thus producing a piecemeal simulation of the apparatus, there is no particular problem in automatically effecting the transfer of the coefficients to the pieces of feedback apparatus. What we have succeeded in doing is to make a white box which can potentially assume the characteristics of any non-linear transducer whatever, and then to draw it into the similitude of a given black-box transducer by subjecting the two to the same random input and connecting the outputs of the structures in the proper manner, so as to arrive at the suitable combination without any intervention on our part.

I ask if this is philosophically very different from what is done when a gene acts as a template to form other molecules of the same gene from an indeterminate mixture of amino and nucleic acids, or when a virus guides into its own form other molecules of the same virus out of the tissues and juices of its host. I do not in the least claim that the details of these processes are the same, but I do claim that they are philosophically very similar phenomena.

This is a section of doi:10.7551/mitpress/11810.001.0001

Cybernetics or Control and Communication in the Animal and the Machine

By: Norbert Wiener

Citation:

Cybernetics or Control and Communication in the Animal and the Machine
By: Norbert Wiener
DOI: 10.7551/mitpress/11810.001.0001
ISBN (electronic): 9780262355902
Publisher: The MIT Press
Published: 2019

Funding for the open access edition was provided by the MIT Libraries Open Monograph Fund.

The MIT Press

First MIT Press paperback edition, February 1965

This book was set in ITC Stone Serif Std and ITC Stone Sans Std by Toppan Best-set Premedia Limited. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Wiener, Norbert, 1894-1964, author.
Title: Cybernetics ; or, Control and communication in the animal and the machine / Norbert Wiener ; forewords by Doug Hill and Sanjoy Mitter.
Other titles: Control and communication in the animal and the machine
Description: [Second edition, 2019 reissue]. | Cambridge, MA : The MIT Press, [2019] | "Reissue of the 1961 second edition." | Includes bibliographical references and index.
Identifiers: LCCN 2019005612 | ISBN 9780262537841 (pbk. : alk. paper)
Subjects: LCSH: Cybernetics. | Control theory. | System theory.
Classification: LCC Q310 .W47 2019 | DDC 003/.5--dc23 LC record available at https://lccn.loc.gov/2019005612

10 9 8 7 6 5 4 3 2 1