

THE MAZE SOLVING COMPUTER

Richard A. Wallace

L. & O. Research and
Development Co.

The maze solving computer is a small digital computing machine mounted on wheels which is able to explore mazes made of toy train track and "learn" the correct path to a predetermined goal. The machine, illustrated in Figure 1, measures 13.6 inches long, 4.6 inches wide, 3.3 inches high, and weighs 7.0 pounds. It has some elemental characteristics of organisms and computers. There are receptors to receive information from guideposts erected alongside the track of a maze, and effectors to provide motion to the machine and to throw the track switches which constitute the choice-points of the maze. The computer's "brain" consists of a director network to make it goal-seek, and a selector-memory network to store and recall past receptor information.

The receptors may be seen in the plan view of the computer shown in Figure 2. They consist of four buttons, Pre-Start, Choice-Point, Dead-

* - A restricted definition of learning suitable for this paper is: "The ability to modify a response to a stimulus because of past experience with the stimulus."

End, and Goal, which operate electrical switches. As the computer moves through a maze, the buttons touch corresponding guideposts mounted beside the track.

The effectors of the computer include the Motor, and the Left and Right Turn Contactors. The motor can be run in forward and reverse directions, and the contactors allow the machine to throw the track switches to the left or the right. There is also an External Start Contactor which is used to start the computer into a maze, and an External Reverse Contactor used to send it back from the goal to the beginning of a maze.

The director network which controls the action of the computer is an assembly of ten miniature relays. The Left Turn Order Relay throws the track switches to the left, while the Right Turn Order Relay throws them to the right. Both relays control the operation of the Forward Control Relay and the Forward Power Relay connected to the motor. A single Reverse Order Relay controls the operation of the Reverse Control Relay and Reverse Power Relay connected to the motor. Motor Trigger Relays





FIGURE 1. THE MAZE SOLVING COMPUTER.






FIG. 2
PLAN VIEW OF MAZE SOLVING COMPUTER

No. 1 and No. 2 are connected together to momentarily interrupt the circuit which self-locks the motor control and power relays. In addition, there is an External Reverse Relay which operates when the computer is returning from the goal to the beginning of a maze.

There are six Memory Switches in the computer which allow it to remember the turns in the path to a goal which may be at most six choice-points removed from the start. A typical Memory Switch is illustrated in Figure 3. It has two coils: a horizontally mounted stepping coil to move the contact arm one step each time it is energized, and a vertically mounted homing coil to return the contact arm to the initial position when it is energized. The contact arm has three positions: the initial or top position initiates a forward and left turn order, the second position a forward and right turn, and the third or bottom position a reverse order.

The Selector Switch illustrated in Figure 4 allows access to and alteration of the Memory Switches. It has two vertically mounted coils: an adder coil to advance the contact arm one step each time it is energized, and a subtractor coil to return the contact arm one step each time it is energized. There are three banks of contacts: a memory storage bank, a memory change bank, and a homing bank. The banks have seven contact positions: an initial position, and an additional position for each Memory Switch. The contacts on the memory storage bank are connected to the contact arms of the Memory Switches, while the contacts on the memory change bank are connected to the coils of the Memory Switches. The homing bank is wired through a pair of normally-closed contact points to the subtractor coil to make it oscillate and return the contact arm to the initial position when the computer reaches the goal of a maze.

Relatively large mazes having one goal and one path to the goal can be solved by the computer. The largest maze which it can solve is shown schematically in Figure 5. It is a branch-like structure having two paths at each choice-point. There are six levels of branches, the number of choice-points on each level increasing in a geometric progression. The

number of levels is called the depth and the number of choice-points at each level is called the breadth of a maze.

The four kinds of guideposts required to actuate the computer's receptor buttons can be located in the scale drawing of a two-level maze shown in Figure 6. A Pre-Start Post is placed at the beginning of the maze. The computer always starts exploring from this post. A Choice-Point Post is placed at the entrance of the maze and before each track switch, and a Dead-End Post is erected at the end of each wrong path. A Goal Post is placed at any track end which the operator designates as the goal.

Four types of contacts are mounted beside the rails of the maze. The External Start Contact is near the Pre-Start Post and touches the corresponding contactor on the computer. Near the Choice-Point Posts are Left Turn and Right Turn Contacts mounted beside one another between




FIG. 3

the rails. They are wired to the solenoids of the adjacent track switch. By energizing them, the computer is able to throw the switch in either direction. The External Reverse Contact is mounted near the Goal Post. When the computer is stopped at the goal, the operator may return it to the pre-start position by energizing this contact. (This operation avoids hand-carrying the machine back to the Pre-Start Post.)

In operation, the computer begins from the Pre-Start Post of a maze after receiving an initial starting pulse from the operator. Figure 7 shows it approaching a choice-point. The computer explores by trial and error, seeking left branches before right, and backing out of dead-end branches until it reaches the Goal Post. It may be thought of as scanning or sampling the branches. At the Goal Post it stops and a lamp on the motor illuminates. The computer can be returned from this position to the Pre-Start Post by the operator. The next time it is started, it will run directly to the Goal Post and stop without having made any wrong choices. Subsequently, should the Goal Post be moved, or should the geometry of the maze be changed, the computer will proceed through the maze in the last set of directions in its memory. If a Dead-End Post is encountered before the Goal Post is reached, the computer will begin exploring and changing its memory until it reaches the new Goal Post. The next time the machine starts into the maze it will proceed directly to the new Goal Post.

Various theories and structures of learning mechanisms have been reported during the past few years. In 1933 Thomas Ross of the University of Washington published an account of a machine which he built. His machine employed a moving arm to explore a maze of parallel passages somewhat resembling a coarse five-tooth comb. In 1938 he built a free-running machine², which could solve deeper mazes than his original moving-arm mechanism. In 1948³, and again in


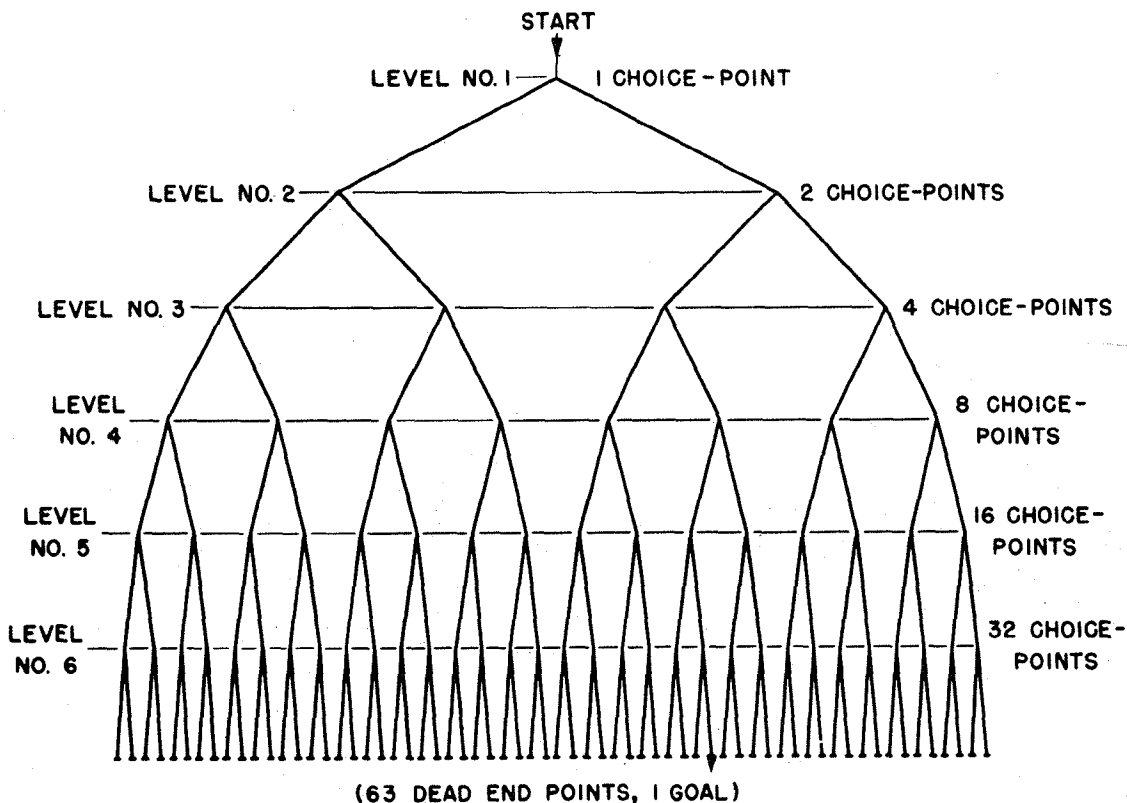


FIG. 4

1950⁴, Norbert Wiener of the Massachusetts Institute of Technology wrote about learning processes, suggested that they could be considered a special sort of feedback, and stated that he thought engineers able to invent mechanisms employing these processes. In 1951⁵, Gray Walter of the Burden Neurological Institute modified his light-seeking machines so they could learn to respond to sound signals, and thereby demonstrated a mechanization of conditioned reflexes. In 1952⁶, Claude Shannon of the Bell Telephone Laboratories completed a maze solving machine having a magnetic "mouse" propelled by a motor running under the maze. The author's interest in learning mechanisms dates from reading Dr. Wiener's books. The present maze solving computer was built to help shed light on the characteristics of learning mechanisms. The machine's behavior appears to



SCHEMATIC DIAGRAM OF LARGEST MAZE COMPUTER CAN SOLVE

FIG. 5

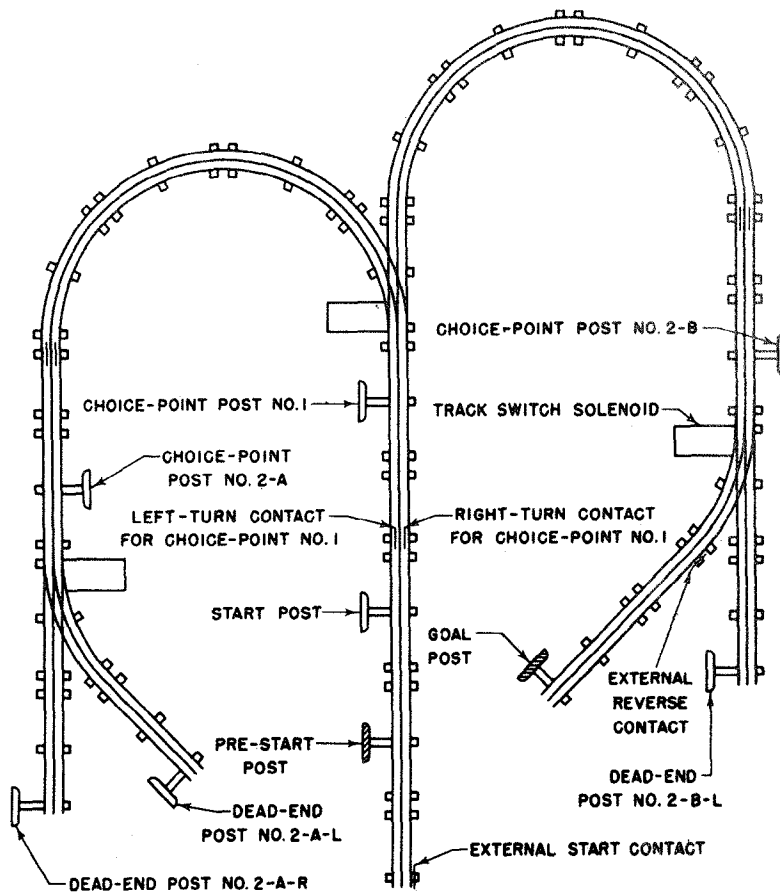


FIG. 6

SCALE LAYOUT OF A MAZE TWO LEVELS DEEP

be explainable largely in terms of present cybernetic theory.

Generalized patterns of information-flow in closed-loop systems are diagramed in Figure 8. The information-flow within a null-seeking mechanism, and between the mechanism and its environment, is illustrated in Figure 8-A; and the flow within a learning mechanism, and between it and its environment, is shown in Figure 8-B.

Environment, as referred to in the paper, may be thought of as a structure capable of producing energy changes which can be sensed or

detected by a mechanism. The nature and resolution of the receptors of the mechanism limit its environment.

Information may be said to flow in the systems of Figure 8 when energy changes occur and cause changes in the entropy, or in the order or disorder of any parts of the systems. Closed-loop information-flow is thought of, here, to include the concept of feedback in which a portion of the output information is fed back into the input. This is in contrast to open-ended information-flow in which no part of the output information reaches an earlier part. The closed-loop flow


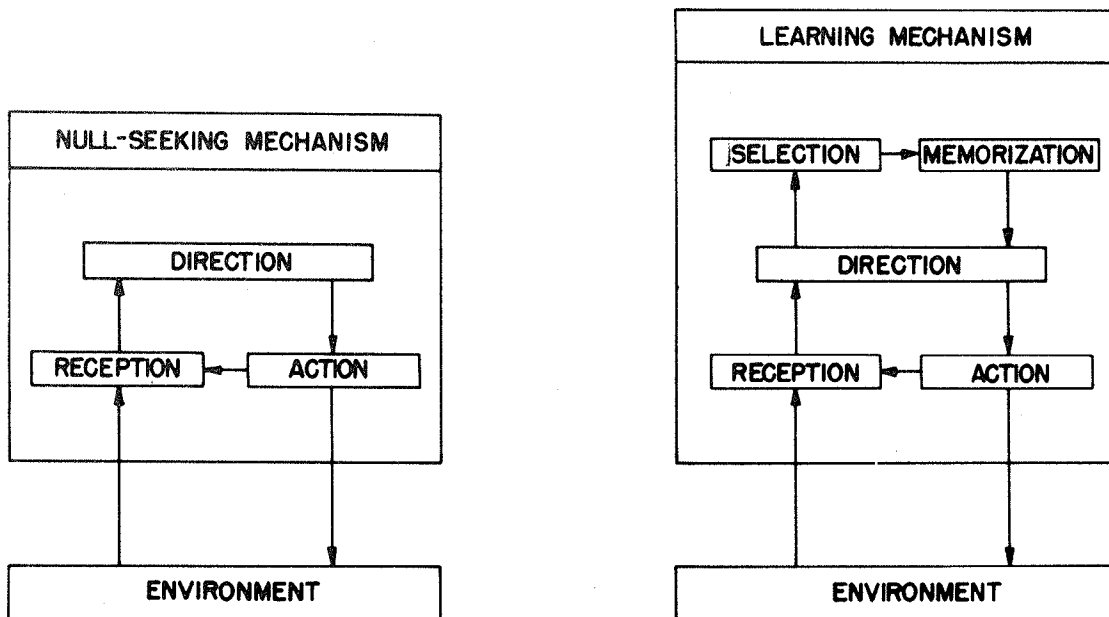


FIG. 7

THE MAZE SOLVING COMPUTER APPROACHING A CHOICE POINT.



A FIG. 8 B

INFORMATION FLOW IN CLOSED-LOOP SYSTEMS

considered is of a negative or regulative nature, as contrasted to positive or oscillatory feedback.

A null-seeking mechanism contains three major functions: direction, reception, and action, as shown in the block diagram of Figure 8-A. The direction block may be thought of as a balancing, regulating, or homeostatic device which endeavors to seek a null-point, or point of minimum potential, stress, or tenseness. It receives information from the reception block, decides whether this information is aiding or hindering the balance it is attempting to main-

tain, and sends out information to the action block of such kind as to cause the desired (balancing) information to be received by the reception block and to reject the unwanted (unbalancing) information. The action block receives information from the direction block and can send it both to the reception block and to the environment. If the mechanism is adapting to its environment, information flows from the action to the reception block; while if the mechanism is altering its environment, information flows from the action block to the environment, and then to the reception block.

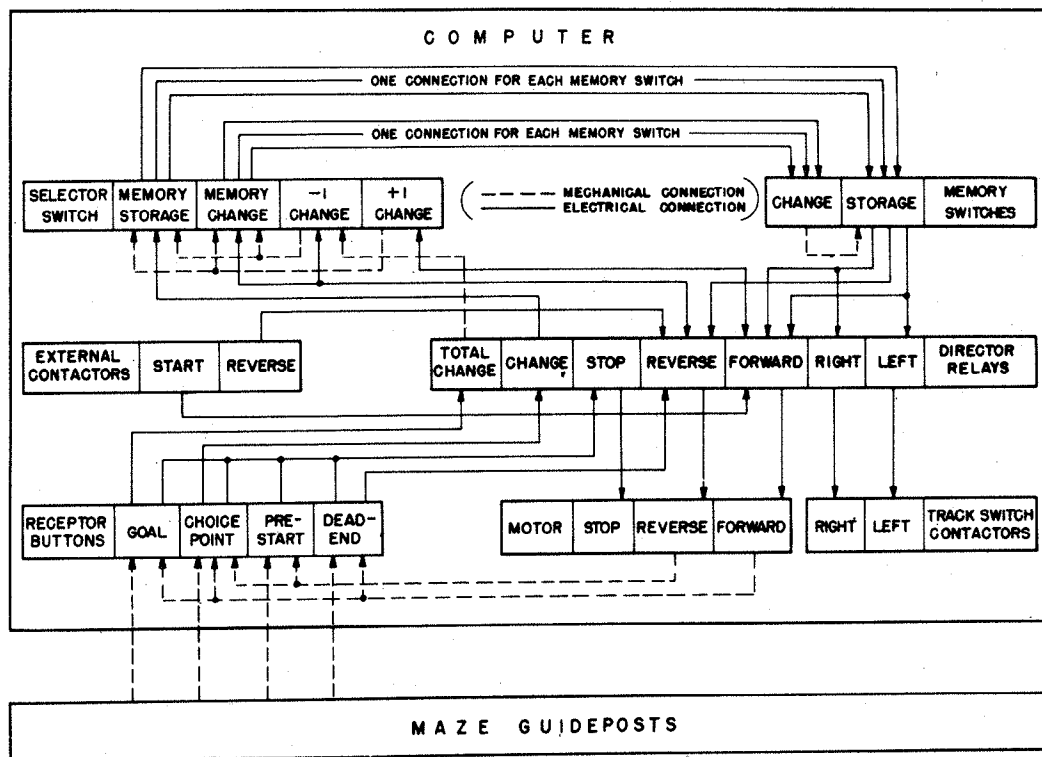


FIG. 9

INFORMATION FLOW OF THE MAZE SOLVING COMPUTER

Thus the minor loop consisting of the reception, direction, and action blocks is called the adapting loop, while the major loop consisting of the reception, direction, and action blocks, plus the environment is called the altering loop. To null-seek, a mechanism must possess either or both of these closed information loops.

A learning mechanism, as shown in Figure 8-B, contains the direction, reception, and action functions of a null-seeking mechanism; and has in addition two more, selection, and memorization. The direction block is the same sort of balancing device found in a null-seeking mechanism, but in addition, it receives information from the memorization block and gives information to the selection block. The selection block which receives information from the direction block and sends it to the memorization block, may be considered the file-clerk of the memory. In addition to having a null-seeking mechanism's adapting and altering loop, a learning mechanism also has a learning loop which connects the direction, selection, and memorization blocks.

The behavior of the direction block in seeking a null-point suggests that the word "problem" can be defined in terms of the internal state of a mechanism rather than in terms of the environment as is more usual. Generalizing, a problem may be said to exist when information sent to the direction block causes its balance to be disturbed. The activity of the mechanism to remove the unbalance is the process of solving the problem. Information which causes the removal of the unbalance within the direction block may be defined as the goal or end which the mechanism is seeking.

It appears that problems, or unbalances of the direction block, are of two general types. One kind occurs when at the beginning of the unbalance, the mechanism can sense the goal with its receptors. This type can be solved by a null-seeking mechanism. The second kind of problem occurs when at the beginning of the unbalance, the receptors of a mechanism can not sense the goal. A maze presents this type problem, and requires a learning mechanism to solve it.

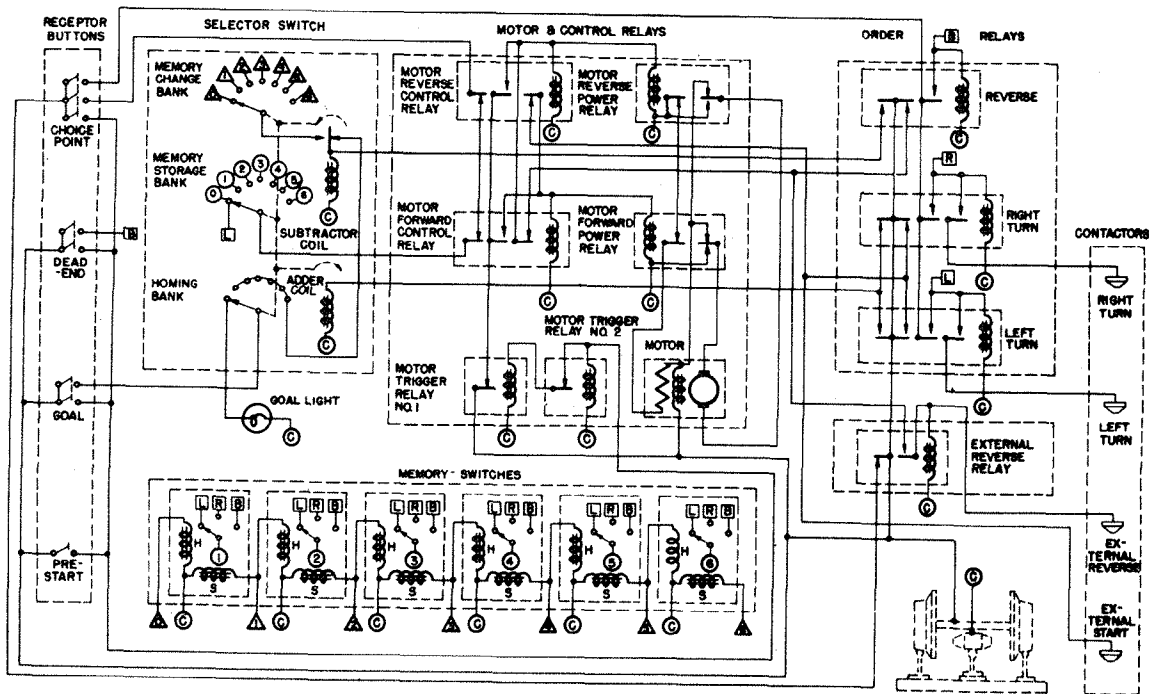
Learning may be thought of as an extension of the senses. A mechanism containing a learning loop can integrate its past with its present receptor information. Thus it is able to solve a class of problems beyond that which it could solve using its receptors alone.

The maze solving computer's information flow is shown in the block diagram of Figure 9. It is similar in form to Figure 8-B except that the major blocks within the computer are subdivided to indicate more completely their functions. The dotted lines indicate mechanical connections, while the solid lines represent electrical connections. The information flow within the computer follows four major paths emanating from the four receptor buttons.

If the Choice-Point Button is depressed, a signal stops the Motor. Another signal is sent through the Director Relays to the Memory Storage contact arm of the Selector Switch. The Memory Storage bank conducts the signal to the Storage contact arm of one of the Memory Switches. The Storage bank sends the signal back to the Director Relays as one of three orders: Left and Forward, Right and Forward, or Reverse. If the order is Left and Forward, the Left Turn Contactor is energized, which throws the track switch left, the Motor is started forward, and the adder coil of the Selector Switch makes a single positive change in the position of the contact arm. If the order from the Memory Switch is Right and Forward, a similar sequence occurs except that the Right Turn Contactor is energized and the track switch is thrown right. Should the order be Reverse, the Motor is started in reverse, the subtractor coil of the Selector Switch makes a single negative change in the position of the contact arm, and a signal is sent to the Memory Change contact arm of the Selector Switch. The Memory Change bank sends the signal to the stepping coil of the appropriate Memory Switch which changes the position of the contact arm.

If the Dead-End Button is depressed, a signal stops the motor. A Reverse signal is sent to the Director Relays which causes the same sequence of changes as if it were sent from the Memory Switches.

FIG. 10



SCHEMATIC DIAGRAM OF THE MAZE SOLVING COMPUTER

If the Goal-Button is depressed, a signal stops the motor. Another signal is sent to the homing tank which causes the Selector Switch contact arm to return to its initial position.

If the Pre-Start Button is depressed, as it is when the computer backs from the Goal Post to the Pre-Start Post, a signal is sent to stop the motor.

The learning and adapting loops of information-flow within the computer can be seen in Figure 9. The ability of the machine to throw the track switches either left or right is not considered to be information-flow going to the maze, because the computer cannot detect the position of the switches and is obliged to go through the operation of throwing them whether necessary or not, each time it runs a maze. The motion of the computer through the maze may be thought of as the mechanical connection between the Motor and the Receptor Buttons. When the computer is returning from the goal to the beginning of a maze, it is operating only as a null-seeking, and not as a learning mechanism.

The actual electrical circuits used in the computer are shown in Figure 10. In addition to the logical circuits required, Figure 10 shows the self-locking circuits and the electrical interlocks used to obtain the desired sequence of operations.

The type and size of the mazes which the present computer can solve is limited. If the machine is placed in a maze which has more than one goal, it will learn the path to the first goal that it finds. If it is placed in a maze having more than one path to the goal it will reach the goal by the first path it discovers--providing, of course, that no runs are deeper than its capacity of six levels.

The maze solving computer is a "one-trip" learner. Any maze which it can solve may be learned in one exploration. The machine can remember only one maze at a time; it learns new mazes at the expense of forgetting past ones. In these respects it differs markedly from ordinary animal and human maze-learning behavior.

The observed differences in behavior of the computer as contrasted with animals and humans can be attributed to a variety of reasons. The computer has been specifically designed to solve a stylized class of mazes, and has been equipped with enough components to accomplish this task. In contrast, animals and humans are much more generalized mechanisms having multiple goals, varying with time, and are equipped with a generous supply of computing elements. The generalized nature of animal and human brains makes possible receiving a variety of information about a maze, much of which is superfluous, and requires additional components to assess statistically the receptor and memory information

to decide what is needed to solve the maze. It seems that the goal of a maze is intermixed with their other goals. To explain this inter-relationship the phenomenon of motivation is introduced. The simplicity of the maze solving computer precludes this complication.

The relays and switches used in the computer are not analogous to animal and human neurons. The connections within the animal and human brain at this time appear so complex that no meaningful comparison with the computer's simple network can be made.

The information acceptance rate of the computer never exceeds the response time of any of its memory components. In contrast, it appears that the speed which an animal or human can move through a maze often exceeds the rate which their brains can recognize and organize the information required for learning the solution.

While the present maze solving computer might be said to demonstrate some of the grosser characteristics of animal and human learning, it is far from any of nature's products. The author hopes that the building of it, and of other allied machines which it has suggested, will help lead to a deeper and broader understanding of the fascinating subject of brain-like mechanisms, and perhaps of the human brain itself.

References:

1. Thomas Ross, "A Machine That Thinks", Scientific American, 1933, Vol. 148, page 206.
2. Thomas Ross, "Synthesis of Intelligence - Its Implications", Psychological Review, 1938, Vol. 45, page 185-9.
3. Norbert Wiener, Cybernetics, John Wiley and Sons, Inc., New York, 1948.
4. Norbert Wiener, The Human Use of Human Beings, Houghton Mifflin Co., Boston, 1950.
5. W. Grey Walter, "A Machine That Learns", Scientific American, August 1951, Vol. 185, No. 2, page 60.
6. John Pfeiffer, "This Mouse is Smarter Than You are", Popular Science, March 1952, page 99.