

READING H-A-S N-O-T D-O-N-E S-O A-L-R-E-A-D-Y.

MACHINE:

JOHN FITCH: You are watching a reading machine, an experimental system which uses a computer programmed to scan a printed page and translate the letters it finds into a series of comprehensible sounds. How this system works and what it may mean as a reading aid for the blind is our story today on *Science Reporter*

[MUSIC PLAYING]

Hello, I'm John Fitch, *MIT Science Reporter*. Today we're at the Research Laboratory of Electronics at MIT to learn about a brand new computer technique for reading printed material out loud. It's a sort of aid that may eventually make such ordinary things as newspapers and personal letters, reading them a possibility for the blind.

Concern for the special problems of the blind is certainly not a recent thing. Way back in the fourth century, AD St. Basil opened the Western world first hospice for the blind at Caesaria Cappadocia in Persia. Then, over the succeeding centuries, such hospices became quite common in Europe, but they all thought only in terms of care, not to education or mechanosensory aids. It was only very slowly that any sort of reading systems were developed.

The first was a scheme for engraving the letters on wooden blocks, though up in 1517 by a Spaniard named Francesco Lucas. Then, over the next few centuries techniques grew to include raised lead type, pinhole marked cardboard, and large wooden letters. All of them ponderous methods, requiring that each copy be made by hand.

The first real breakthrough came in the late 18th century when a Frenchman named Valentin Hauy, who was director of an institute in Paris, discovered that one of his pupils could decipher several letters on a card because of the indentations made by the printing. Hauy immediately used his pen handle to emboss some more signs which the boy could also read. The result was printing in relief.

Once the principle of relief printing became commonplace, the well-known Braille alphabet was adapted for use throughout the English speaking world. Now, Braille typewriters and stencils permit the blind to communicate in the written medium.

As useful is Braille is, however, it still has some severe limitations. To learn about what some of these problems are and what modern technology can do about them, we talked with Dr. Samuel Mason, a professor of electrical engineering at MIT, and leader of a cognitive information processing research group at the Research Laboratory of Electronics.

SAMUEL MASON: Well, as you know John, there are quite a few people in this research group who are interested in problems of sensory aids for the deaf, blind, and the blind. I think most people know about the cane that blind people use to help them avoid obstacles, and no doubt, many people heard of Braille, the touch language that allows blind people to read printed material when it's been transcribed into Braille. There are also these talking books, pre-recorded by human voice so that blind people can listen to novels and general literature.

But, these aids that exist, and certain other technical systems that have been developed, don't really solve all problems of the blind person. What we're going to be looking at today is an experimental reading machine system that will solve some of the other problems.

For instance, if a blind person gets his morning newspaper, he can't wait for it to be transcribed into Braille. You have to send it somewhere and get it back. And then there's the specialized material, mathematics and so on, that maybe never will be put in Braille form. Then there's reading matter that needs to be read immediately, personal correspondence or assistive material such as bills and dunning letters. It's really towards the immediate reading problem that we're directing ourselves.

I'd like to show you a machine, an experimental system, that we have which will have a spelled speech output. It'll recognize print and speak out the letters. If it sees the word cat it'll say C-A-T on the output. Let's go over and take a look at it.

I think we'll first look at the front end here. Here's the thing called the carriage and we've cut a piece out of a page a book. This is an honest sample of book print. We're not cheating, we didn't have the specially made up.

JOHN FITCH: That's not much of a plot.

SAMUEL MASON: No. We want to look at realistic material. We've arranged it this way for purposes of testing the system. Of course, in its final form, the machine would not require the user to cut out strips of paper from a book.

JOHN FITCH: It would read the whole page.

SAMUEL MASON: Yes, it would work on the book itself. So, let me put this in here. There we are. I'll move it over to there. Now underneath this box is a thing called a flying spot scanner. This is very much like a home TV set. There's a screen here with a bright spot it that can be moved around to by the rest of the equipment.

There's a lens in here, John, and that bright spot is focused up right on the surface of that printed of material and there's a very tiny bright spot there. If it falls on the white part of the page a lot of light is reflected, picked up by a photo cell and we know that we're on white. If it falls on black, there's not much like reflected and we know we're on black. So, the machine could tell the difference between white and black, which is all you need to get the information off the page.

JOHN FITCH: Can we see it work?

SAMUEL MASON: Yes. I'm going to cover this up because there's quite a bit of light in here and this little box is not completely light tight. So, let me turn on the scanner. The Switch is 9 and 11. Come on over here a second, John.

This is a small digital computer that we built here to control that scanner.

JOHN FITCH: It'll tell where to move the spot around?

SAMUEL MASON: It tells the spot where move, depending on what it's seeing. Actually, we're testing out various systems. In final form we wouldn't need such complexity, but we do need this for flexibility to see what to do.

I'm going to ask Dr. Lee, Dr. Francis Lee, who was one of the two project engineers on this system development to run the scanner for us.

JOHN FITCH: All right. Now what will we be seeing?

SAMUEL The first thing we'll see here is just sort of a TV display of the particular word that's on the scanner.

MASON:

JOHN FITCH: Oh, yes, I can see the word away.

SAMUEL Yeah, the spot is moving across the whole field there, just the way it does on your TV set. But, what we want is to

MASON: tell what shape the letter is.

JOHN FITCH: Otherwise this isn't very useful, is it?

SAMUEL This particular information doesn't do us much good. Francis can now switch into another mode of operation, in
MASON: which the little spot, under the control this special purpose computer, actually makes decisions as it goes along. It senses whether it's on black or white and traces around the edge of the letter.

JOHN FITCH: Now we have just the letter W.

SAMUEL Its sort of like the little spot is like a man walking around an island, maintaining his path by telling whether he's
MASON: walking in water or on sand. He makes a decision and keeps on the line. So, this spot can trace around the outline the letter.

JOHN FITCH: I don't see quite how a spot makes up his mind whether he's on the black or on the white.

SAMUEL We'll have to look at some details if we want to really see how it works. So, let's go over to the blackboard, John. I
MASON: have cheated just a little bit by putting spots on the blackboard. What these are are the specific locations where that scanner looks. It doesn't look everywhere, it just looks at certain spots.

I've made a very much expanded picture here. These spots are really only about 0.02 of an inch apart, or maybe even closer. I'll take an example. If the scanner happened to have a capital D on it than the black region would be in there. If I want to blow up the upper left hand corner, here, I could say that maybe the D is coming in like this.

JOHN FITCH: I see, it really is expanding.

SAMUEL Yeah, well it really should be expanded even more. So, in here is black, that's all black area over here. There's
MASON: the inside corner here, but we don't have to worry about that. Now, let's say that the scanner gets going, I want to start it up here, I think. One, two, three, yeah that's what I want to do.

Suppose it gets going from a previous point and here is senses white. That's outside the letter. The computer is programmed to tell the spot to tell the scanner to turn right, make a 90 degree right turn when it stops and senses white. So, it jumps to the right, then, it makes a right turn and goes over to here.

Now it stops and says what color is? And it's black. So, now it's instructed to make a left turn. Left, turn for black and a right turn right for white. It makes a left turn, goes up here, and now it sees black again, so it makes another left turn. And here we go, right turn in white, right turn.

Now it's going to do something a little trickier. It's got to make three successive left turns because it saw three black points. It's back to the white point again so it turns right and comes out. Now, as you can see from this expanded picture, it looks a little bit like garbage but remember that here's the edge of the letter. We really have these spots very close together, so the thing is looking around here and really drawing us a pretty good picture of the edge of the letter.

JOHN FITCH: I didn't see a wiggly line on that screen. What was that I was seeing?

SAMUEL
MASON: The thing with displayed on the screen was only the black points. This one, this one, this one, this one, this one, this one, this one.

JOHN FITCH: They really are along the edge.

SAMUEL
MASON: So, if you connect all of them, that's the thing you saw. The points we're so close together look like a continuous line, but it was really these separate points. Now, this gives us the edge points, but that doesn't recognize the letter, so we have to go further and say what the computer does when it's handling this information.

Here's how it goes, and I'm going to pick a different, letter just for convenience, because it's a nice one to look at. That's the letter C. We actually do a little smoothing in a few other steps here that I'm going to leave out, but we wind up with this edge information, this edge trace, on these successive points. Now, what we do, as the computer marches through this contour it is instructed to pick out the extremes of the emotion in north-south and east-west.

Here's obviously the place where you get farthest north on this letter, and we'll make a mark there. Now, we come on over here and we get to a place where we're way over to the east, so we'll make a mark there. Now, here comes down to the south, you make a mark there. That isn't the most southerly point, but it is a place where it turns from south to north, so we mark that, also. Here's a north and west and a south and then a north and then an east and a south and back here to a west, again,

JOHN FITCH: Now, when you say mark it, you don't mean actually make little lines in the picture. Does it sort of remember where these are?

SAMUEL
MASON: No, what the computer does is just store the information on which of these events happened in what order. And, really, what we store in the computer is a 0 for an east or west and a 1 for north or south. So, this would be 0,1, 0, 1, 1, 0, 1, 1, 0, 1.

Now, this scanner actually sneaks in from the left and finds the letter, then it starts around. If it starts out here and would go 0, this is all that would be retained in the computer, plus a little bit of other stuff that I won't go into. A 0 and then 1,0, 1, 1, 0, 1, 1. That is called a code word.

JOHN FITCH: Oh, I see. That describes that particular shape.

SAMUEL
MASON: Yeah, and as a matter of fact, the fellow who was originally working on this got so he could look at these words pretty much tell what the shape looked like.

JOHN FITCH: Yes, but that doesn't say to the computer that that's the letter C. How does that 0101--

SAMUEL Well, we have to train the computer. You see, what we've done here is we've come down from a whole lot of
MASON: points to just the points on a line, and then we selected just a few points off of that line. So, we're getting rid of all this extra information. We're getting it down to simpler information, but still retaining enough so that we can tell which letter it is.

What happens, as we'll see in a moment, have to train a computer. You have to let it look at a letter and then the computer tells you that it picked up a certain code word, but it says to you, I don't know what it is. You tell it, you type on the typewriter over there, C. From then on it knows that this is C, it puts a C by it.

JOHN FITCH: Could we see that work?

SAMUEL Yeah, we keep going that way.

MASON:

JOHN FITCH: All Right.

SAMUEL Professor Donald Troxel, the other co-project engineer on this development, is here at the console and he's going
MASON: to do a little teaching and training of the machine for us. I think what we'll do first is to set up a word that the scanners looking at. Without any stored memory all, starting fresh, the machine now knows nothing, we'll start tracing letters and generating these code words that I mentioned. Then we'll have to tell the machine what each word means, which letter it designates. So, let's try the first one.

They're going through the preliminary set up mode. They've traced the letter A and generated a code word here 401452536.

JOHN FITCH: Those aren't just 1s and 0s there.

SAMUEL Well, that's a different kind of coding. That's called an octal coding and each of those numbers there represents a
MASON: group of three 1s and 0s.

JOHN FITCH: Oh, I see. It sort of adds them up.

SAMUEL If it were 1, 1, 1 it would write 7. It's binary coding. Now Don can train the machine by typing the letter A. The
MASON: machine now knows that that's an A. Now it's traced another letter, the next one, W, in that word away. And we'll start training it on the W. And now it catches another A and generates different code word.

JOHN FITCH: Why didn't it know that was A? We taught it that already, didn't we?

SAMUEL Well, when it sees a different A in the first place, the printed shape of that might be a little bit different due to ink
MASON: run, but not too much different.

JOHN FITCH: So it said 6536 instead of 2536.

SAMUEL The only difference, here, is in this first number in the second part of the code group.

MASON:

JOHN FITCH: OK, so we could tell it that's an A, also.

SAMUEL Yeah. Now comes a Y, the letter Y, and we'll train it on that and that's the word away. Shall we go back and see
MASON: how it does?

JOHN FITCH: OK, what did it write? Oh, period. It was a period at the end of it. It already knows periods, huh?

SAMUEL It knows periods because if there's a little fly speck that isn't too small, then that's a period.
MASON:

JOHN FITCH: All right. Could we are now try the word again and see if it knows it's this time?

SAMUEL OK, let's do that. Don has return it and it shouldn't-- it missed that A again because it generated different code
MASON: word. We've got A, W and Y, period. What happened here? It missed the A again. Let's try once more. Can we try that again? Here we go on that word. OK, new code word, we tell it it's an A, new code word for the W, and it got the W-A-Y period.

So, it isn't doing too well, yet. It's sort of stupid, hasn't learned enough.

JOHN FITCH: I could understand why it might miss two different As because imperfections in the printing, but why would a miss the same letter like W, which is only there once.

SAMUEL For instance, that first letter in the word away. When the scanner sees that letter the second time, remember that
MASON: this is a real scanner working on real printed matter, and there's a little shimmer, a little jitter, the shape might not look exactly the same to the scanner the second time. Things don't happen exactly the same with a real live Skinner, so it is possible for the letter to generate several different code words even when looking at the same character.

In fact, if we wanted to train this up in a hurry on one letter, we could just to turn off the scanner advance and pick one of these letters and let it look at it several times and see how long it takes to store enough code words to put that letter in the bag.

Here we go again. A-W, OK. W,W,W, it's doing pretty well. Oops, let's stop. Here at the got the A, then it got the W. It retraced the W about a dozen times and it recognized it each time. Which meant that each time there was generated, on that same W, in that same position one of the code words that we already have for W. Then it finally generated a new code word that it hadn't seen before because of this jitter. So, we'll tell it that that's a W.

Now, we'll see how it does. Oh, we picked up another one.

JOHN FITCH: Now it knows about four different code words for W.

SAMUEL Here comes a fifth one then. Now it's doing pretty well, W, W, W, W, W.
MASON:

JOHN FITCH: Each time around it says W.

SAMUEL Yup. Well, let's let it go for a second or two. Found another one.
MASON:

JOHN FITCH: So, what you're saying is you might really need to store away about, oh, half a dozen different code words or same letter.

SAMUEL Yeah, we actually can do a trick when we're running the machine and trying to run it usefully with is some stored
MASON: memory-- that's enough-- when we're trying to run it with some stored memory, we have say, five or six of these
code words stored, if we trace a letter and it generates a new code word we can tell it no, SOAP, go back and
trace it again. Then chances are, we'll get one of the word stored in our memory. In fact, we make a trace as
many as two or three times if necessary, and that way the machine does pretty well with a minimum of storage.

JOHN FITCH: Haven't you already trained this machine so it can read text?

SAMUEL Yeah, we put in some time with the material from this book that we're training it on and we have a tape which
MASON: we're now actually running into the machine over in the other building, the computer that we're tied to, that does
this recognition. It's a PDP over in building 26, and we're running the tape in now and setting up the full memory
that we've already accomplished by training.

JOHN FITCH: You've done some technique like this of teaching how read W.

SAMUEL Yes.
MASON:

JOHN FITCH: You've got few codes for that and a few for A.

SAMUEL Yup, somebody sat here and trained it on some material and trained it up on each letter like this. I think we have
MASON: about 400 or 500 stored code words and this should be enough to take care of this alphabet pretty well.

We're going to look at a new line now. They're setting up the computer over there. Is the tape reader on?

MAN 1: You already read the tape in.

MAN 2: Thank you.

SAMUEL OK, that's our intercom to the other building. Don is starting it on the beginning of the line instruction. Shall I
MASON: read off the line?

JOHN FITCH: Let's see how the computer makes out with the line.

SAMUEL It's sort of finding the beginning of the line, now. Now what happened? And he started to lead.
MASON:

JOHN FITCH: And he started to lead. That's pretty good.

SAMUEL OK, and he started to lead, that's a portion of the text. It looks to me as though we had a little drop the pulse or
MASON: hung up there somewhere at the beginning. Should we try once more?

JOHN FITCH: Sure.

SAMUEL Has not so already, period.
MASON:

JOHN FITCH: It still doesn't like the capital A.

SAMUEL No. And he started to lead. All right, now how did it do? Let's see how it did. The portion of the material it was
MASON: looking at was the following: has not done so already, period. And he started to lead. The only character it missed was a capital A.

JOHN FITCH: Well, that's very interesting, but tell me Dr. Mason, this is wonderful for it to be able to recognize, but if it's just types it out again, that isn't going to be very useful for the blind person.

SAMUEL No, he can't do much better with the typewritten page than the original book page. So, we'll go over to the
MASON: speech. But , just before we do that, I to mention that we've done about 50 characters here and have got one of them wrong, but we have done some testing on this material with larger blocks of material and it's doing about half a percent error.

So, out of every 200 characters it might miss one, but this isn't too bad for a reading machine for the blind. In fact, some people think that a blind person could operate with a reading machine with an error rate as high as 5%. We're down to 1/10 of that right now.

Now let's see what the speech sounds like.

JOHN FITCH: All right.

READING H-A-S.
MACHINE:

SAMUEL Has.
MASON:

READING N-O-T.
MACHINE:

SAMUEL Not.
MASON:

READING D-O-N-E.
MACHINE:

SAMUEL Done. Next word is all ready.
MASON:

READING A-L-R-E-A-D-Y, P-E-R-I-O-D.
MACHINE:

JOHN FITCH: Oh, it spells out the period.

READING A-N-D.
MACHINE:

JOHN FITCH: And.

READING H-E.
MACHINE:

JOHN FITCH: He.

READING MACHINE: S-T-A-R-T-E-D.

SAMUEL MASON: Started.

READING MACHINE: T-O.

SAMUEL MASON: To.

READING MACHINE: L-E-A-D.

JOHN FITCH: Lead.

SAMUEL MASON: Incidentally, this time it got that letter it missed before. The capital A.

JOHN FITCH: Very interesting, but if you were trying to read a book, would be an awfully slow and laborious means of doing it, wouldn't it?

SAMUEL MASON: Yeah, well, the way we have this set up now, we're reading the spelled letters out slowly so that it's easy to understand what's going on. But, these letters were prepared from human speech, that is a human being-- Ken Bingham, actually, who did this work-- spoke into a microphone and then his speech wave was sampled and stored on tape as numbers, and it's in the machine over there, and he has tailored these sounds down until each one only takes 2/5 of a second. Actually, only one fifth of the second, 200 milliseconds.

So, if we packed this close together right now it would be, instead of saying already, A-L, it would be A-L-R-E-A-D-Y. Fast, at about 60 words a minute.

JOHN FITCH: Right.

SAMUEL MASON: Now, we're doing some research work on compressing that further by overlapping the edges of these sounds. We have some tape that's been pre-prepared on the computer to illustrate that compression.

JOHN FITCH: Tell me, is there any possibility that in the future you might actually be able to generate real words, not spell them out?

SAMUEL MASON: Yeah, that's another phase of this project. We are working on connected speech. With connected speech, instead of saying A-L-R-E-A-D-Y, I would say already. This is more of a job to do.

JOHN FITCH: Would you store the syllables themselves somehow?

SAMUEL The first part of that would require translating from the English spelling into what's called a phonetic spelling,
MASON: because the phonetic symbols are what you need to drive an artificial speech generator. Now, we think we know how to do this. We hope we can accomplish that within a year or two. That would be, of course, at a much higher rate, you could maybe make that go 180 words a minute.

JOHN FITCH: What other sorts of things working on?

SAMUEL For our studies of the human information requirements and capabilities of blind people we really want to get a big
MASON: system together, here, that'll allow us to see what kinds of things a blind person can learn to do and I can find it comfortable to do. So, we have a lot of flexibility here. We're working on other outputs such as a tactical display, a sort of a slow speed tactical TV, which can lay out pictures and diagrams and other material that the blind person couldn't get from a recognition system.

JOHN FITCH: So you'd actually feel the outlines of a picture. Thank you, very much, Dr. Mason. It's most interesting.

Today we've been visiting the Research Laboratory of Electronics at MIT. Our guest has been Dr. Samuel Mason.
I'm John Fitch, *MIT Science Reporter*

[MUSIC PLAYING]