

REITH LECTURES 1984: Minds, Brains and Science

John Searle

Lecture 2: Beer Cans & Meat Machines

TRANSMISSION: 14 November 1984 – Radio 4

In my last lecture, I provided at least the outlines of a solution to the so-called ‘mind-body problem’. Mental processes are caused by the behaviour of elements of the brain. At the same time, they’re realised in the structure that’s made up of those elements. Now, I think this answer is consistent with standard biological approaches to biological phenomena. However, it’s very much a minority point of view. The prevailing view in philosophy, psychology and artificial intelligence is one which emphasises the analogies between the functioning of the human brain and the functioning of digital computers. According to the most extreme version of this view, the brain is just a digital computer and the mind is just a computer program. One might summarise this view—I call it ‘strong artificial intelligence’, or ‘strong AI’—by saying that the mind is to the brain as the program is to the computer hardware.

This view has the consequence that there’s nothing essentially biological about the human mind. The brain just happens to be one of an indefinitely large number of different kinds of hardware computers that could sustain the programs which make up human intelligence. On this view, any physical system whatever that had the right program with the right inputs and outputs would have a mind in exactly the same sense that you and I have minds. So, for example, if you made a computer out of old beer cans powered by windmills, if it had the right program. It would have to have a mind. And the point is not that for all we know it might have thoughts and feelings, but, rather, it must have thoughts and feelings, because that is all there is to having thoughts and feelings—implementing the right program.

Now, most people who hold this view think we have not yet designed programs which are minds. But there is pretty much general agreement among them that it’s only a matter of time until computer scientists and workers in artificial intelligence design the appropriate hardware and programs which will be the equivalent of human brains and minds.

Many people outside the field of artificial intelligence are quite amazed to discover that anybody could believe such a view as this. So, before criticising it, let me give you a few examples of the things the people in this field have actually said. Herbert Simon of Carnegie Mellon University says that we already have machines that can literally think. Well, fancy that! Philosophers have been worried for centuries about whether or not a machine could think, and now we discover that they’ve already got such machines at Carnegie-Mellon. Simon’s colleague Alan Newell claims that we have now discovered (and notice that Newell says ‘discovered’ and not just ‘hypothesised’ or ‘considered the possibility’, but actually discovered) that intelligence is just a matter of physical symbol manipulation: it has no essential connection with any kind of biological or physical wetware or hardware. Both Simon

and Newell, to their credit, emphasise that there's nothing metaphorical about these claims: they mean them quite literally.

Marvin Minsky of MIT says that the next generation of computers will be so intelligent that we will 'be lucky if they're willing to keep us around the house as household pets'. But my all-time favourite in the literature of exaggerated claims on behalf of the digital computer is from John McCarthy, the inventor of the term 'artificial intelligence'. McCarthy says even a machine as simple as a thermostat can be said to have beliefs. And, indeed, almost any machine capable of problem-solving can be said to have beliefs.' I admire McCarthy's courage. I once asked him: 'What beliefs does your thermostat have?' And he said: 'My thermostat has three beliefs—it believes it's too hot in here, it's too cold in here and it's just right in here.' Now, as a philosopher, I like all these claims for a simple reason. Unlike most philosophical theses, they are reasonably clear, and they can be simply and decisively refuted. It's this refutation that I am going to undertake in this lecture.

The nature of the refutation has nothing whatever to do with any particular stage of computer technology. I think it's important to emphasise that, because in these discussions the temptation is always to think that the solution to our problems must wait on some as yet uncreated technological wonder. But, in fact, the nature of the refutation is completely independent of any state of technology. It has to do with the very definition of a digital computer, with what a digital computer is.

It's essential to our conception of a digital computer that its operations can be specified purely formally; that is, we specify the steps in the operation of the computer in terms of abstract symbols—sequences of zeros and ones printed on a tape, for example. A typical computer 'rule' will determine that when a machine is in a certain state and it has a certain symbol on its tape, then it will perform a certain operation such as erasing the symbol or printing another symbol and then enter another state such as moving the tape one square to the left. But the symbols have no meaning; they have no semantic content; they're not about anything. They have to be specified purely in terms of their formal or syntactical structure. The zeros and ones are just numerals: they don't even stand for numbers.

Indeed, it's this feature of digital computers that makes them so powerful. One and the same type of hardware, if it's appropriately designed, can be used to run an indefinite range of different programmes. And one and the same program can be run on an indefinite range of different types of hardware. But this feature of programs, that they're defined purely formally or syntactically, is fatal to the view that mental processes and program processes are identical. And the reason can be stated quite simply. There's more to having a mind than having formal or syntactical processes. Our internal mental states, by definition, have certain sorts of contents. If I'm thinking about Kansas City or wishing that I had a cold beer to drink or wondering if there will be a fall in interest rates, in each case my mental state has a certain mental content in addition to whatever formal features it might have. That is, even if my thoughts occur to me in strings of symbols there must be more to the thought than the abstract strings, because strings themselves can't have any meaning. If my thoughts are to be about anything, then the strings must have a meaning which makes the thoughts about those things. In a word, the mind has more than a syntax, it has semantics,

Now, the reason that no computer program can ever be a mind is simply that a computer program is only syntactical, and minds are more than syntactical. Minds are semantical, in the sense that they have more than a formal structure, they have a content.

To illustrate this point, I've designed a certain thought-experiment. Imagine that a bunch of computer programmers have written a program that will enable a computer to simulate the understanding of Chinese. So, for example, if the computer is given a question in Chinese, it will match the question against its memory, or data base, and produce appropriate answers to the question in Chinese. Suppose, for the sake of argument that the computer's answers are as good as those of a native Chinese speaker. Now then, does the computer on the basis of this, understand Chinese, does it literally understand Chinese, in the way that Chinese speakers understand Chinese?

Well, imagine that you are locked in a room, and in this room are several baskets full of Chinese symbols. Imagine that you (like me) do not understand a word of Chinese, but that you are given a rule-hook in English for manipulating these Chinese symbols. The rules specify the manipulations of the symbols purely formally, in terms of their syntax, not their semantics. So a rule might say: 'Take a squiggle-squiggle sign out of basket number one and put it next to a squoggle-squoggle sign from basket number two. Now suppose that some other Chinese symbols are passed into the room, and that you are given further rules for passing back Chinese symbols out of the room. Suppose that, unknown to you, the symbols passed into the room are called 'questions' by the people outside the room, and the symbols you pass back out of the room are called 'answers to the questions'. Suppose, furthermore, that the programmers are so good at designing the program, and that you are so good at manipulating the symbols, that very soon your answers are indistinguishable from those of a native Chinese speaker. There you are, locked in your room, shuffling your Chinese symbols and passing out Chinese symbols in response to incoming Chinese symbols. On the basis of the situation as I have described it, there's no way you could learn any Chinese simply by manipulating these formal symbols.

The point of this whole story is simply this: by virtue of implementing a formal computer program, you behave exactly as if you understood Chinese, but all the same, you don't understand a word of Chinese. But now, if going through the appropriate computer program for understanding Chinese is not enough to give you an understanding of Chinese, then it's not enough to give any other digital computer an understanding of Chinese. And again, the reason for this can be stated quite simply, if you don't understand Chinese then no other computer could understand Chinese because no digital computer, just by virtue of running a program, has anything that you don't have. All that the computer has is, as you have, a formal program for manipulating un-interpreted Chinese symbols. To repeat, a computer has a syntax, but no semantics.

The whole point of the parable of the Chinese room is to remind us of a fact that we knew all along. Understanding a language, or, indeed, having mental states at all, involves more than having just formal symbols. It involves having an interpretation, or a meaning attached to those symbols. And a digital computer, as defined, cannot have more than just formal symbols because the operation of the computer, as I said

earlier, can only be defined in terms of its ability to implement programs. And these programs are purely formally specifiable—that is, they have no semantic content.

I think we can see the force of this argument if we contrast what it's like to be asked and to answer questions in English, and to be asked and to answer questions in some language where we have no knowledge of any of the meanings of the words. Imagine that in the Chinese room you are also given questions in English about such things as your age, or your Life-history, and that you answer these questions. Now, what's the difference between the Chinese case and the English case? Well, again, if, like me, you understand no Chinese and you do understand English, then the difference is obvious. You understand the questions because they are expressed in symbols whose meanings are known to you. Similarly, when you give the answers in English you are producing symbols which are meaningful to you. But in the case of the Chinese, you have none of that. In the case of the Chinese, you simply manipulate formal symbols according to a computer program, and you attach no meaning to any of the elements.

Now, various replies have been suggested to this argument by workers in artificial intelligence and in psychology, as well as philosophy. They all have something in common; they are all inadequate. And there's an obvious reason why they have to be inadequate, since the argument rests on a very simple logical truth, namely, syntax alone is not sufficient for semantics, and digital computers, in so far as they are computers have, by definition, a syntax alone.

I want to make this clear by considering a couple of the arguments that are often presented against me. Some people attempt to answer the Chinese-room example by saying that the whole system understands Chinese. The idea here is that though I, the person in the room manipulating the symbols, do not understand Chinese, I'm just the central processing unit of the computer system. And they argue that it's the whole system, including the room, the baskets full of symbols, the ledgers containing the programs and, perhaps other items as well, taken as a totality, that understands Chinese. But this is subject to exactly the same objection I made before. There's no way that the system can get from the syntax to the semantics. I, as the central processing unit have no way of figuring out what any of these symbols means, but then neither does the whole system.

Another common response is to imagine that we put the Chinese-understanding program inside a robot. If the robot moved around and interacted causally with the world, wouldn't that be enough to guarantee that it understood Chinese? Once again,, the inexorability of the semantics-syntax distinction overcomes this manoeuvre. As long as we suppose that the robot has only got a computer for a brain, then even though it might behave exactly as if it understood Chinese it would still have no way of getting from the syntax to the semantics of Chinese. You can see this if you imagine, once again, that I am the computer. Inside a room in the robot's skull I shuffle symbols without knowing that some of them come in to me from television cameras attached to the robot's head and others go out to move the robot's arms and legs. As long as all I have is a formal computer program. I have no way of attaching any meaning to any of the symbols.

So suppose the robot picks up a hamburger and this triggers the symbol for hamburger to come into the room. Well, as long as all I have is the symbol, with no knowledge of

its causes or how it got there, I have no way of knowing what it means. The causal interactions between the robot and the rest of the world are irrelevant unless those causal interactions are represented in some mind or other. But there's no way they can be if all that the so-called mind consists of is a set of purely formal, syntactical operations.

It's important to see exactly what is claimed and what is not claimed by my argument. Suppose we ask the question that I mentioned at the beginning: 'Could a machine think?' Well, in one sense, of course, we're all machines. We can construe the stuff inside our heads as a meat machine. And, of course, we can all think. So, trivially, there are machines that can think. But that wasn't the question that really bothered us. So let's try a different formulation of it.

Could an artefact think? Could a man-made machine think? Well, once again, it depends on the kind of artefact. Suppose we designed a machine that was, molecule for molecule, indistinguishable from a human being. Well then, if you can duplicate the causes, you can presumably duplicate the effects. So, once again, the answer to that question is, in principle at least, trivially yes. If you could build a machine that had the same structure as a human being, then presumably that machine would be able to think. Indeed, it would be a surrogate human being. Well, let's try again.

The question isn't 'Can a machine think?' or 'Can an artefact think?' The question is 'Can a digital computer think?' But now, once again, we have to be very careful in how we interpret the question. From a mathematical point of view, anything whatever can be described as if it were a digital computer. And that's because it can be described as instantiating or implementing a computer program. So in an utterly trivial sense, the pen that's on the desk in front of me can be described as a digital computer. It just happens to have a very boring computer program. The program says: 'Stay there.'

Now, since in this sense anything whatever is a digital computer because anything whatever can be described as implementing a computer program, then, once again, our question gets a trivial answer. Of course our brains are digital computers, since they implement any number of computer programs. And of course our brains can think. So, once again, there's a trivial answer to the question. But that wasn't really the question we wanted to ask. The question we wanted to ask is this: 'Can a digital computer, as defined, think?' That is to say, is instantiating or implementing the right computer program with the right inputs and outputs, sufficient for, or constitutive of, thinking? And the answer to this question, unlike its predecessors, is clearly 'no'. And it's 'no' for the reason that we have spelt out, namely, the computer program is defined purely syntactically. But thinking is more than just a matter of manipulating meaningless symbols, it involves meaningful semantic contents. These semantic contents are what we mean by meaning'.

It's important to emphasise again that we're not talking about a particular stage of computer technology. My argument has nothing to do with the forthcoming, amazing advances in computer science. No doubt we will be much better able to simulate human behaviour on computers than we can at present, and certainly much better than we have been able to in the past. But the point I am making is that if we are talking about having mental states, having a mind, all of these simulations are simply

irrelevant. It doesn't matter how good the technology is, or how rapid the calculations made by the computer are. If it's really a computer, its operations have to be defined syntactically, whereas consciousness, thoughts, feelings, emotions and all the rest of it involve more than a syntax. Those features, by definition, the computer is unable to duplicate, however powerful may be its ability to simulate. The key distinction here is between duplication and simulation. And no simulation by itself ever constitutes duplication.

What I've done so far is to give a basis to the sense that those quotations I began this talk with are really as preposterous as they seem. But there's a puzzling question in this discussion, and that is, 'Why would anybody ever have thought that computers could think or have feelings and emotions and all the rest of it?' After all, we can do computer simulations of any process at all that can be given a formal description. So, we can do a computer simulation of the flow of money in the British economy, or the pattern of power-distribution in the Labour Party. We can do computer simulations of rain storms in the Home Counties, or warehouse fires in East London. Now, in each of these cases, nobody supposes that the computer simulation is actually the real thing; no one supposes that a computer simulation of a storm will leave us all wet, or a computer simulation of a fire is likely to burn the house down. Why on earth would anyone in his right mind suppose a computer simulation of mental processes actually had mental processes? I don't really know the answer to that, since the idea seems to me, to put it frankly, quite crazy from the start. But I have a couple of speculations.

First of all, where the mind is concerned, a lot of people are still tempted to some sort of behaviourism. They think if a system behaves as if it understood Chinese, then it must really understand Chinese. But we've already refuted this form of behaviourism with the Chinese-room argument.

Another assumption made by many people is that the mind is not part of the biological world; it's not part of the world of nature. The strong artificial intelligence view relies on that in its conception that the mind is purely formal; that somehow or other, it cannot be treated as a concrete product of biological processes like any other biological product. There is in these discussions, in short, a kind of residual dualism. AI partisans believe that the mind is more than a part of the natural biological world; they believe that the mind is purely formally specifiable. And the paradox of this is that the AI literature is filled with fulminations against some view called dualism, but in fact the whole thesis of strong AI rests on a kind of dualism. It rests on a rejection of the idea that the mind is just a natural biological phenomenon in the world like any other.

I want to conclude this lecture by putting together the thesis of the last lecture and the thesis of this one. Both of these theses can be stated very simply. And indeed, I'm going to state them with perhaps excessive crudeness. But if we put them together, I think we get a quite powerful conception of the relations of minds, brains and computers. And the argument has a very simple logical structure, so you can see whether it's valid or invalid.

1. Brains cause minds.

Well, of course, that's really too crude. What we mean by that is that mental processes that we consider to constitute a mind are caused, that is, entirely caused, by processes

going on inside the brain. But let's be crude; let's just write that down as three words—brains cause minds. And that's just a fact about how the world works. Now let's write proposition number two:

2. Syntax is not sufficient for semantics.

That proposition is a conceptual truth. It just articulates our distinction between the notions of what is purely formal and what has content. Now, to these two propositions - that brains cause minds and that syntax is not sufficient for semantics - let's add a third and a fourth:

3. Computer programs are entirely defined by their formal, or syntactical, structure.

That proposition, I take it, is true by definition; it's part of what we mean by the notion of a computer program. Now let's add proposition four:

4. Minds have mental contents; specifically, they have semantic contents.

And that, I take it, is just an obvious fact about how our minds work. Now, from these four premises, we can draw our first conclusion; and it follows obviously from two, three, and four; namely:

This is conclusion 1: No computer program by itself is sufficient to give a system a mind. Programs, in short, are not minds, and they are not by themselves sufficient for having minds.

That's a very powerful conclusion, because it means that the project of trying to create minds solely by designing programs is doomed from the start. This is a purely formal, or logical, result from a set of axioms which are agreed to by all (or nearly all) of the disputants concerned. That is, even the most hardcore enthusiast for artificial intelligence agrees that in fact, as a matter of biology, brain processes cause mental states, and they agree that programs are defined purely syntactically. But if you put these conclusions together with certain other things that we know, then it follows immediately that the project of strong AI is incapable of fulfilment.

However, once we've got these axioms, let's see what else we can derive. Here's a second conclusion:

Conclusion 2: The way that brain functions cause minds cannot be solely in virtue of running a computer program.

And this second conclusion follows from the first premise, as well as from our first conclusion. That is from the fact that brains cause minds and that programs are not enough to do the job, it follows that the way that brains cause minds can't be solely by running a computer program. Now that, also, I think is an important result because it has the consequence that the brain is not, or at least is not just, a digital computer. We saw earlier that anything can be trivially described as if it were a digital computer, and brains are no exception. But the importance of this conclusion is that the computational properties of the brain are simply not enough to explain its functioning to produce mental states. And indeed, that ought to seem a commonsense scientific conclusion to us anyway, because all it does is remind us of the fact that brains are biological engines; that biology matters. It's not, as several people in artificial

intelligence have claimed, it's not just an irrelevant fact about the mind that it happens to be realised in human brains.

Now, from our first premise, we can also derive a third conclusion:

Conclusion 3: Anything else that caused minds would have to have causal powers at least equivalent to those of the brain.

And this third conclusion is a trivial consequence of our first premise: it's a bit like saying that if my petrol engine drives my car at 75 miles per hour, then any diesel engine that was capable of doing that would have to have a power output at least equivalent to that of my petrol engine. Of course, some other system might cause mental processes using entirely different chemical or biochemical features from those that the brain in fact uses. It might turn out that there are beings on other planets or in some other solar system that have mental states and yet use an entirely different biochemistry from ours. Suppose that Martians arrived on earth, and we concluded that they had mental states. But suppose that when their heads were opened up, it was discovered that all they had inside was green slime. Well, still, that green slime, if it functioned to produce consciousness and all the rest of their mental life, would have to have causal powers equal to those of the human brain. But now, from our first conclusion, that programs are not enough, and our third conclusion, that any other system would have to have causal powers equal to the brain, conclusion four follows immediately:

Conclusion 4: For any artefact that we might build which had mental states equivalent to human mental states, the implementation of a computer program would not by itself be sufficient. Rather the artefact would have to have powers equivalent to the powers of the human brain.

The upshot of this entire discussion I believe is to remind us of something that we've known all along: namely, mental states are biological phenomena. Consciousness, intentionality, subjectivity, and mental causation are all a part of our biological life-history, along with growth, reproduction, the secretion of bile and digestion.