## What are computers (if they're not thinking things)?

Book or Report Section

Accepted Version

Preston, J. ORCID: https://orcid.org/0000-0002-3740-2308 (2012) What are computers (if they're not thinking things)? In: Cooper, S. B., Dawar, A. and Lowe, B. (eds.) How the World Computes: Turing Centenary Conference and 8th Conference on Computability in Europe, CIE 2012, Cambridge UK June 2012 Proceedings. Lecture Notes in Computer Science, 7318. Springer-Verlag, Berlin and Heidelberg, pp. 609-615. ISBN 9783642308697 doi: https://doi.org/10.1007/978-3-642-308703_61 Available at https://centaur.reading.ac.uk/28691/

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.
Published version at: http://www.springerlink.com/content/u251k144r0135145/
To link to this article DOI: http://dx.doi.org/10.1007/978-3-642-30870-3_61
Publisher: Springer-Verlag
Publisher statement: The original publication is available at www.springerlink.com
All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the End User Agreement.

## www.reading.ac.uk/centaur

## CentAUR

Central Archive at the University of Reading
Reading's research outputs online

## What are Computers (if they're not Thinking Things)?

Many of us now imagine that in the future humans, either will, or at least could, 'in theory', construct an electronic digital computer which would really be a thinking thing. Alan Turing was perhaps the most notable exponent of this view, and a significant proportion of his published work was devoted to arguing for $i t$.

Even if one accepts this 'computationalist' view, the question 'What are computers?' is still worth asking, since almost no-one thinks of past and existing computers, or even of most foreseeable computers, as thinking things. We still need an account of what these devices that now surround us are, even if computationalists are right to think of them as proto-thinkers, as it were. I think that when one sees the answer to this question, the temptation to think that more sophisticated computers really would be thinking things evaporates.

## Our Technologies

The electronic digital computer is the most prominent among a proliferation of devices whose operations we regularly and quite naturally describe in some terms that can also be used to characterise human actions: mousetraps catch mice, washing machines wash clothes, dishwashers wash the crockery, thermostats turn the central heating on and off, pocket calculators calculate, guided missiles seek the exhaust heat of aircraft in order to destroy them, etc, etc.

Sometimes the underlying idea is still that these devices are things that we use in order to carry out these tasks, activities or functions. In the case of some of these devices, though, we have somehow learned to take seriously our talk of the device itself carrying out the activity. Usually, no misunderstanding results, but the computer falls into this group, since we think and talk of computers as if
they compute, calculate, search, execute instructions, etc. This is entirely contingent, incidentally: when Turing was writing, the term 'computer' meant person who performs computations, and it could easily have kept that meaning. If someone had invented a handy term or acronym for what we now call computers (e.g. 'EDMs', for 'electronic digital machines'), 'computer' could have kept its original meaning, and we would easily and naturally have thought of EDMs as 'the devices that computers use'.

There are three aspects to this kind of attribution: the task gets done, it gets done correctly, and it's the device that does it. To take all this at face value would be to conclude that these devices can correctly be said to be actually performing that function or carrying out that activity which would otherwise have to be performed or carried out by a person, using the activity-term in the same sense that it has when applied to humans. So we can speak of computers not only as 'searching for' data, 'storing' information, 'executing' instructions, but also as playing chess, scheduling tasks, controlling the operation of other machines, etc.

I'd like to suggest, though, that this way of talking when applied to computers encourages us to conceive of them in the wrong way, that it involves getting their entire relation to us wrong. In the case of computers, it's quite remarkable that hardly anyone stands up for this alternative and, I think, commonsense view of them and their activities. A relatively superficial feature of the language we use to talk about such devices is allowing computationalists to set the terms of the debate about 'machine intelligence'.

## Replacing-Technologies

We human beings aren't the only tool-using animals, but by any standard our use of tools is the most widespread and the most impressive. We use tools in situations where we want something done but don't want to use (only) our
bodies to do it. Our reasons for not using only our bodies can be various: not wanting to get hurt, not being able to fit into or reach a particular space, not wanting to get dirty, performing some task (like driving in a screw, or cutting a piece of wood) for which no part of the human body will do, as well as considerations of speed, cost, efficiency, etc. Alongside using tools (in this core sense) that relieve us of tasks that are too dangerous, dirty, or physically unmanageable for our bodies, we now also use technologies that relieve us of tasks that are boring, complicated, or expensive. This is often where computers come in. Of course, we also use computers to help us control machines of other kinds - cars, telephones, industrial machines, etc..

Computationalists are right to think that there's something importantly different about the computer. The computer is a new kind of technology. But it's a kind that's nevertheless contiguous with other kinds of technologies, and to understand it as a technology is to understand both this contiguity and the way in which it goes beyond previous technologies.

For each such technology there's a crucial distinction between the process or activity involved and the product that results. What I call 'replacingtechnologies' are machines and devices that replace human activities. Not all our machines are replacing-technologies: no number of human beings, unaided by technology, could produce certain electromagnetic phenomena, or atomic power, or a nuclear explosion. We use replacing-technologies to ensure that the product or end-state of some activity that could only formerly be done by humans is brought about, even though the process or activity itself hasn't taken place. Of course, some process has taken place, but this process isn't the one that we formerly used, unaided, to produce that result.

Our pre-computational technology mostly involved tools and implements for the performance of tasks that used to require a significant amount of human physical effort. Computers, though, mostly relieve us of a different kind of activity, the kind of activity that used to require thinking, reasoning, calculating,
and certain other psychological skills. Because the activities they replace are of this kind, intellectual rather than manual, and because these activities are often quite distinctive to human beings (they can't be performed by other creatures), computers are a very special, and especially impressive, kind of replacementtechnology, a technology capable of being especially close to our minds (and thus perhaps to our hearts, as it were).

This distinctiveness is of course down to Alan Turing, who showed not only that a certain kind of mathematical function can be encoded in a way that makes it ideally suited to be mechanically implemented (in electrical circuitry), but also that what we now call 'Turing machines' are an absolutely general kind of device which can execute any procedure that consists in the manipulation of symbols. Since almost all our programmed electronic devices can be thought of as (horrendously complicated) Turing machines, our stored-program digital computers are general-purpose technologies in a specific way, a way in which no previous technology has ever been. Their ever-increasing importance to us derives at least partly from this fact.

## Intentional Actions and Non-Intentional Operations

Because the main philosophical debate has been about machine intelligence, computationalists tend to think that if one could show that intelligent operations can be 'broken down into' unintelligent ones, and then show that Turing machines operate entirely on the basis of millions of operations of just this unintelligent kind, one would have successfully opened the door to the idea of an intelligent machine. There are problems with this strategy, but regardless of whether it works, I suggest instead that the key issue is really not intelligence but another, more fundamental feature of certain psychological concepts, the one that philosophers call intentionality. Intentionality is a feature of certain psychological phenomena - crucially, the same psychological phenomena that computationalists are centrally concerned
with, such as thoughts, beliefs, and desires. Its technical definition and history aren't vital here but, the thing to note is that whether or not the high-level operations of computers are rightly described in 'intentional' terms (as computationalists would like), those operations can always be broken down into sub-operations that are non-intentional. At the lowest level, (the level of the Turing machine), the operations involved are clearly and purely mechanical. They are operations such as: scanning a square on the machine-tape, registering the contents of that square, erasing the symbol in the square, writing such a symbol, etc. The fact that these operations are purely mechanical is a sure sign that they're not genuinely intentional. Neither human actions nor human psychological skills, though, can be thus 'broken down'. Complex actions and skills can be thought of as composed of simple ones, of course, but these simple ones are still intentional. This is a logical feature of intentionality, not merely a contingent one, and it marks a fundamental, categorical divide between the way in which computers work and the nature of human beings. However complex our computers become, they will only involve 'more of the same' (i.e. more such mechanical operations), and this can't yield thought or intelligence.

## The Intentional Stance?

Isn't it the case, though, as Daniel Dennett has suggested, that we take the intentional stance with respect to computers? That is, that we (at least on occasion) think of them in intentional terms, ascribing intentional properties to them?

It's true that we do occasionally say of computers that they 'want' or are 'waiting for' an instruction, that they're 'recalling' certain information, or even (when they're taking a long time processing) that they're 'thinking about' an issue. More significantly, as I've already stressed, although adult human beings don't often think of non-computational devices in such psychological terms, we do describe computers as if they achieve the tasks they in fact replace.

But to conclude from either of these habits that we do or can take the 'intentional stance' towards them would be a wild exaggeration. The range of intentional terms that we apply to computers is radically restricted and thus a very weak analogue of the full range of intentional terms that can be applied to humans. There's simply no application for most of the associated psychological terms we predicate of other humans, such as hoping, fearing, expecting, anticipating, regretting, considering, reconsidering, contemplating, pondering, reflecting, etc. This is a consequence of what philosophers call the 'holism' of the intentional: the serious application of an intentional term brings with it the actuality, or at least the possibility, of applying other intentional terms to that same person. What's more, even when we consider those few intentional terms that we do apply to computers, we should remember that humans are of course quite capable of applying them to non-computational devices, and even to artefacts which aren't devices at all (such as dolls). So although computers (machines generally) can be described using a very weak image or shadow of one range of our psychological vocabulary, it does no harm whatsoever to think that our applications of such terms to machines, including computers, are always in 'scare-quotes', and not to be taken seriously.

Nevertheless, to talk of past, present or, I would argue, future computers as achieving the tasks they replace, or even occasionally as performing intelligent activities, such as thought, does at least have two very good rationales. The first is the fact that what their activities replace are indeed usually stretches of thought or intellectual activity. Where this clearly isn't the case, as for example in the case of a car's engine management system, which is also computational, there's no temptation to think of the device in intentional terms. So where we humans, unaided, would have to think in order to attain some result, a temptation to conceive also of the computer which can be made to achieve that result (perhaps faster, more reliably than we can) as a thinking thing is quite naturally felt.

The second rationale for talking of computers using a very limited range of intentional terms is that most of us simply don't know how computers operate. Their operation is recondite, that is, hidden from our inspection, and its full mechanical description is simply unwieldy. It would be quite wrong though, to suggest, as Dennett sometimes does, that the only way of coming to understand 'what the computer is doing' is from the 'intentional stance'. That is, we all know that there is a purely mechanical description of their activity, although most of us couldn't say what it is. Nevertheless, taking the 'intentional stance' towards computers (where this is understood in a very limited way, as a matter of predicting and explaining their operations and products in terms of goals, searches, decisions, etc.) is quite natural when one knows of no other way of explaining their operations, or even when this other way is simply too longwinded (which is almost always the case with computers, simply due to the number of operations they are capable of performing at great speed). This also explains why computer science texts usually explain the operation of computers in terms of analogies with human rule-following activities (such as games, recipes, knitting patterns, etc.).

Both these rationales, however, are very clearly pragmatic in nature. They can do nothing to support the serious philosophical conclusion that such machines are really performing intentional operations. It's surely ironic that serious thinkers who would otherwise scout such considerations have allowed them to prevail in this context, since elsewhere such rationales would normally be regarded as hopelessly insufficient for a philosophical conclusion.

## Conclusion

In sum, computers are the very latest kind of labour-saving technologies. The labour they save isn't that of the 'workers by hand', though, but that of the 'workers by brain'. However sophisticated, computers aren't things which compute, but things that we use to replace the human activity of computation.

Perhaps they might thus be described as things that people compute with, or things that people use to compute. But even these descriptions might have to be taken with a pinch of salt.

Among those working in the field of artificial intelligence, it is perhaps the designers of 'expert systems' who have come closest to realising the replacement status of computer technologies, since they're quite explicitly aware that they are (in their terms) trying to 'mechanise' the boring bits of expert's jobs. How ironic that the least interesting bit of artificial intelligence (from the point of view of philosophy) should have the clearest conception, not merely of the goals of the subject, but also of the nature of the devices it produces!

Where does all this leave Alan Turing himself? Turing manifestly still counts as a great thinker, one whose mathematical work made possible the development of devices that are effecting an enormous and unprecedented transformation in human work- (and leisure-) activity. He is thus the herald of a new machine-age (for better or worse, of course). But if we cease to think of him as the herald of thinking or intelligent machines, his achievements will in no way be underestimated.

John Preston,
Department of Philosophy, The University of Reading

