



© 1997–2004, Millennium Mathematics Project, University of Cambridge.

Permission is granted to print and copy this page on paper for non–commercial use. For other uses, including electronic redistribution, please contact us.

---

January 2001

Reviews

## 'Computers, Ltd.'

reviewed by Mark Wainwright



## Computers Ltd.: what they *really* can't do

By David Harel

Reviewed by [Mark Wainwright](#)

Computers can do many things, but there are some things they can't do. They certainly can't play tennis or the violin, but those aren't the kinds of thing we're concerned with. There are computational questions, questions of the kind that we would naturally turn to a computer to help us with, that, in fact, they cannot answer (and nor, therefore, can we).

There are two kinds of unanswerable question. Some kinds of questions cannot be answered by any computational procedure at all, under any circumstances. A well–known example is the question: "Will this computer program, given this data, produce an answer – any answer, no matter whether it be the right one – or will it just hang the computer for ever?" This is the so–called *Halting problem*, and the fact is that it cannot itself be answered by a computer program, however much processing power it is given, however long you are prepared to wait for an answer, and however clever the programmer.

The second kind of unanswerable question is the kind that can be computed, but the computation simply takes too long – so long that even if, say, the computer were the size of the known universe, it would take billions of years to finish the calculation. Not just because our technique is deficient, since we can prove that any program at all will take this long; and not because we require it to operate on a mind–bogglingly large amount of data, since the input data can be rather small.

These startling facts have been well known for many years; results about non-computability go back to the late 1930s, for example. They are not merely theoretical, but affect a whole host of real applications. So it is surprising that they have not had much popular coverage, particularly when the closely-related, but somewhat more abstract, incompleteness theorems of Goedel have had plenty of exposure.

## Pretty pictures

All this has changed with the appearance of David Harel's lovely little book. In four chapters, he presents with meticulous clarity the fundamental results of computation and complexity theory to a non-mathematical audience. The fifth chapter briefly surveys some experimental models of computing (such as quantum computing and molecular computing), to see whether these might, in the future, allow us to paint a different picture. In the sixth he takes the opportunity, having talked about intractable problems, to explain how their very intractability can be turned to good use in the creation of what appear to be truly uncrackable codes – the now enormously important field of public-key cryptography. A final chapter waxes a little more philosophical on the difficulty of creating artificial intelligence.

Several things make this book appealing. One is its physical appearance: it is unthreateningly small in size, and well laid out with plenty of pictures. Harel has a knack for the apt and illuminating analogy, so the exposition is never in danger of becoming too abstract. His occasional footnotes are a model of how footnotes should be used – to concisely nail some misapprehension before it arises in the reader's mind, without distracting from the thrust of the argument. (Though, speaking of footnotes, a kindly editor should have told the author what a twit he looks trying to "correct" an often "misquoted" folk rhyme.) And Harel's own research interests happily include problems that not only well illustrate his theme, but lend themselves to attractive diagrams.

A particular merit of the book is that it does not easily allow the reader to walk off with the wrong end of the stick. The results Harel presents are surprising if one has not seen them before, and it is easy to suppose that they are statements about what is possible with some limited view of computing, or with today's technology, rather than fundamental limits on the power of computers. Harel repeatedly returns to this point, and leaves his readers very clear about their true significance.

A perhaps inevitable consequence of this firmness is that he is inclined to overstate his case. For example, readers who accept his frequent assurances that in real life, "exponential-time" algorithms are invariably impracticable will be surprised to learn, in a later aside, that the simplex algorithm is one such and has been happily in use for many years. But this is a fairly venial transgression. The repeated vehemence with which we are told that these intractability results are for all computers, for all time, is another matter. For example, Harel asserts in the introduction that the book

concentrates on bad news that is *proven*, *lasting* and *robust*, concerning problems that computers are simply not able to solve, regardless of our hardware, software, talents or patience. And when we say "proven", we mean really proven; that is, mathematically, and not just experimentally.

*Mathematically*, as Harel knows, quantum computing offers the possibility of fast solutions to problems that are otherwise truly intractable (though not, as he notes, to incomputable ones). In a short section on the subject, he concentrates on the practical problems of building a quantum computer. Without doubt these are formidable, but so were those of building any kind of digital computer in, say, Babbage's time; not all that long ago on the scale of things. Perhaps Harel has simply underestimated quantum computing, and in any event rhetoric like the above serves a valuable purpose as I noted before, but all the same he ought to be a little more careful with it.

## Artificial intelligence

Harel's final chapter is a bit of a departure from the rest of the book. In it, he tries to convince us that machines will never display "human" intelligence, in the sense, for example, of Turing's famous "imitation game" (the "Turing test") where a computer is programmed to answer the questions of a human interrogator in another room in such a way as to convince the interrogator that it is in fact another human. The supposed impossibility of such a program has nothing whatever to do with the mathematically-proven impossibilities of earlier chapters. It's unfortunate that, on this point, Harel is at far less pains than before to prevent the reader getting the wrong impression. After so much emphasis on mathematically proven limitations, it would have been worth stressing that these don't and can't underlie the difficulty of creating "artificial intelligence" (AI), since we humans exhibit the real thing with a brain whose individual neurons could be adequately modelled by a digital computer.

Like the rest of the book, the chapter is entertaining and well written; but fundamentally it seems out of place. Its real message is that David Harel does not believe so-called Artificial Intelligence (AI) is going to happen any time soon, or probably ever. That is interesting in itself, but it is a fact about David Harel, not about computers, and the chapter is lacking in anything much like a reasoned argument. The closest it comes is when Harel urges the reader, "Try to think of the way we humans would have dealt with the situations [discussed], and of the hopelessness of programming a computer to deal with them intelligently."

This sounds like something between a failure of imagination and a pious hope. Fifty years ago, in the original 1950 paper from *Mind* where he first proposed the Turing Test, Alan Turing showed that he fully appreciated how difficult was the test he was proposing, and yet that he was even then quite capable of imagining a machine that would pass it. (Anyone interested enough to have read this far should read the paper; it is sparkingly readable, entirely non-technical, and still provides one of the most cogent arguments on any side of the AI debate.) Turing describes a number of objections that people are inclined to raise to the possibility of AI. One, which he calls the "Heads in the Sand" objection, he characterises thus: "The consequences of machines thinking would be too dreadful. Let us hope and believe that they cannot do so." It is, he says, particularly prevalent among intellectual people, since they most value the power of thinking. Turing notes that this position needs no refuting: "Consolation would be more appropriate." Today the Heads in the Sand objection is alive and well. Its high priest is Roger Penrose, and its bible his book *The Emperor's New Mind*. David Harel, it seems, is one of its acolytes.

### In brief

*Computers Ltd* is a beautiful, well-written, enjoyable book. It presents with extreme clarity results which are important and which ought to be better known. If the final chapter is more speculative, that is no crime – I might have thought it a virtue if I had happened to agree with it. Anyone who is interested in what we know about computing should read it.

#### **Book details:**

Computers Ltd.: what they *really* can't do  
David Harel  
Hardback – 240 pages (2000)  
Oxford University Press  
ISBN: 0198505558

---

'Computers, Ltd.'



*Plus* is part of the family of activities in the Millennium Mathematics Project, which also includes the NRICH and MOTIVATE sites.