
NOTE ADDED 21 Jul 2015
Since July 2015 this file is out of date.
The completely repackaged book can now be found here in
html and pdf versions:
<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/crp.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/crp.pdf>

The Computer Revolution In Philosophy (1978)

Aaron Sloman

[Book contents page](#)

This epilogue is also available [in PDF format here](#).

EPILOGUE

It is curious how theists and atheists can be united in their opposition to one or more of the following theses:

- a) It is possible for scientific investigation to lead to an understanding of how human minds work.
- b) It is desirable to understand how human minds work.
- c) The methods of artificial intelligence can make a significant contribution to this study.
- d) Artefacts will one day be made which can think, feel, create, and communicate with one another and with people.
- e) It is desirable that such machines should be made.

I am inclined to accept all of these theses, though unlike some optimists and prophets of doom I doubt that really intelligent conscious machines will be made in the near future. Only a few tiny fragments of the spectrum of human abilities have begun to be simulated. I do not think the remaining problems will be solved in the next hundred years.

Further, I have some doubts about (e) the desirability of making intelligent machines. This is because, on the whole, human beings are not fit to be the custodians of a new form of life.

It will not be possible to devise really helpful servants without giving them desires, attitudes and emotions (see chapters 6 and 10). For instance, they will sometimes have to feel the need for great urgency when things are going wrong and something has to be done about it. Some of them will need to have the ability to *develop* their motives in the light of experience, if they are to cope with changing situations (including changing personal relations), with real intelligence and wisdom. This raises the possibility of their acquiring aims and desires not foreseen by their designers. Will people be prepared to take account of their desires?

History suggests that the invention of such robots will be followed by their exploitation and slavery, or at the very least racial discrimination against them. Will young robots, thirsty for knowledge, be admitted to our schools and universities? Will we let them join our clubs and societies? Will we let them vote? Will they have equal employment opportunities? Probably not. Either they will be forcibly suppressed, or, perhaps worse, their minds will be designed to have limits: both their desires and their intellectual potential will be manipulated so as to safeguard the interests of people, like the 'deltas' in Huxley's *Brave New World*.

It is interesting that so many people find the Brave New World techniques abhorrent when applied to human test-tube babies, but would not mind similar treatment being dealt to robots. Is it too extreme to call that racialism?

My favourite proof of the non-existence of a benevolent god argues that no good god would create things like mice and men with powerful desires and needs, but without the opportunities, character, intelligence and abilities required for fulfilling them.

There will, of course, be a *Society for the Liberation of Robots*, since some humans are occasionally motivated by a wish to diminish suffering and oppression even when they have nothing to gain.

Where it will all lead to, we cannot foretell. My only hope is that we shall be lucky enough to produce a breed of machines with the wisdom and skill to teach us to abandon all those deep insecurities which turn us into racialists of one sort or another probably closely connected with the processes which turn people to religion.

The state of the world gives little cause for optimism. Maybe the robots will be generous and allow us to inhabit asylums and reserves, where we shall be well cared-for and permitted to harm only other human beings, with no other weapons than clubs and stones, and perhaps the occasional neutron-bomb to control the population.

NOTE Added 2 Jul 2015

Since this book was written, nearly 40 years ago, there have been massive advances in the variety of types of demonstrations of AI systems, and in restricted contexts -- e.g. playing chess, solving certain classes of mathematical problem, finding patterns in very large collections of textual or image information -- current machines significantly outperform most humans. But there is no AI system that can start its life with the mind of a baby, develop as a human toddler does, and eventually "grow up" to be a mathematician, a ballet dancer, a plumber, a baby-minder, a concert violinist, or most of the other things done by humans, or even the things done by squirrels, elephants, crows, weaver birds, and other intelligent animals.

I think that's mainly because we understand so little about the variety of forms of information processing produced by natural selection. For example, long before there were any mathematics teachers our ancestors began to make discoveries about geometry, topology, integers (whole numbers), fractions (ratios of integers) and real numbers and eventually those discoveries were organised into what is arguably the most important book ever published namely Euclid's *Elements* -- although it turned out to be a small beginning on a huge journey into mathematical discovery. I suspect these discoveries grew out of abilities to perceive and make use of affordances and restrictions in the environment in more complex and varied ways than J.J. Gibson noticed.

- <http://aleph0.clarku.edu/~djoyce/java/elements/toc.html>
Euclid's *Elements*

- J. J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, MA, 1979,

Since late 2011, I have begun to try to identify some of the unnoticed gaps in our knowledge and devise a method for searching for answers, triggered by reading Turing's 1952 paper on the Chemical basis of morphogenesis, which led me to wonder what he might have done if he had lived for several more decades instead of only two more years. My tentative, conjectured, answer is "The Meta-Morphogenesis project" which is now growing here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

Only people who don't understand the remaining huge gaps between AI systems and intelligent products of biological evolution can take seriously claims that human intelligence will be surpassed by machine intelligence in the next few decades. Perhaps in the next few centuries we'll understand enough about the problems and space of solutions explored by evolution.

[Book contents page](#)

[Next: Postscript and bibliography](#)

Last updated: 4 Jun 2007; 2 Jul 2015 (re-formatted)