
Huge but unnoticed gaps between current AI and natural intelligence

Aaron Sloman

<http://www.cs.bham.ac.uk/~axs/>

[School of Computer Science, University of Birmingham](#)

Paper accepted for PTAI17 Leeds University 4-5 Nov 2017

<https://www.pt-ai.org/2017>

<https://www.pt-ai.org/2017/programme>

(Revised 11 Oct 2017)

(html file plus derived pdf file)

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-ptai17.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-ptai17.pdf>

Abstract

Despite the huge practical importance of developments in AI, there have always been researchers (including Alan Turing) less interested in using AI systems to do useful things and more interested in potential of AI as **science** and **philosophy**; in particular the potential to advance knowledge by providing new explanations of natural intelligence and new answers to ancient philosophical questions about what minds are. Particularly deep questions ask how biological evolution produced so many different forms of intelligence -- a diverse subset of the space of possible minds, including humans and non-human animals, and humans at different stages of development, in a huge variety of physical and cultural contexts. We don't seem to be close to discovering how to build machines that can replicate all known forms of natural intelligence. I am not claiming that the task is impossible, though the education of AI researchers (and many others) blinds them to some of the important natural phenomena that current AI cannot model (some of them discussed by Immanuel Kant about 240 years ago). In particular, current AI systems that I know of are not even close to matching the amazing discoveries by ancient mathematicians, including discoveries that remain in widespread use all over this planet by scientists, mathematicians, engineers, architects and others. There are deep, mostly unnoticed, connections between ancient discoveries in geometry and topology, and the intelligence of many non-human animals and pre-verbal human toddlers. The discovery processes required are unlike statistical/probabilistic learning, for reasons spelled out in Kant's philosophy of mathematics. Perhaps recognising these limitations will inspire more researchers to join a search for extensions to current AI mechanisms.

See notes: [\[1\]](#), [\[2\]](#).

Towards a more encompassing set of mechanisms

The practical uses of AI, and the rate at which they are now multiplying have been so impressive that some serious thinkers have begun to fear that we are in danger of building monsters that will take over the planet and do various kinds of harm to humans, that we may be unable to prevent because we don't match their intelligence. For some reason most such thinkers don't consider the

more optimistic possibility, suggested many years ago[3], that truly superhuman intelligence will include a kind of wisdom that rejects the selfish, thoughtless, competitive, destructive, superstitious, and other objectionable features that lead to so much harm done by humans to other humans and other species. But "singularity risks" are not my concern now: this paper is about how [little](#) progress has been made in philosophical and scientific aspects of AI that motivated the early researchers who hoped, as I still do, that AI can give us powerful new ways of modelling and understanding natural intelligence: AI as *science* and *philosophy* not just *engineering*.

Alas, AI as engineering dominates AI education (and publicity) nowadays, in contrast with the concerns of early researchers in the field, including some philosophers, who noticed the potential of research in AI to contribute to our understanding of natural intelligence. (Examples: [Turing\(1950\)](#), [Miller et al.\(1960\)](#), [Simon\(1967\)](#) [Minsky\(1968\)](#), [Clowes\(1967\)](#), [McCarthy&Hayes\(1969\)](#), [Boden\(1978\)](#), [Sloman\(1978b\)](#), [Dennett\(1978\)](#), [Dennett\(1996\)](#), and [McCarthy\(1996/2008\)](#). For a deep and wide-ranging survey starting centuries earlier, see [Boden\(2006\)](#).)

Recent spectacular engineering successes mask deep limitations in scientific and philosophical progress in AI. Two results of this masking (at present) are a shortage of good researchers focusing on the unsolved problems, and a shortage of funds for long term scientific research. A European Commission initiative announced in 2003, funded from 2004, [Maloney \(2003\)](#) temporarily shifted the focus of robotics research in the EU back to science, but (as happened in the UK Alvey Project begun two decades earlier), the initiative expanded too rapidly at a time when too few people had had the right sort of education, and also demanded practical demonstrations far too soon. As a result the focus changed in later EU projects to demonstrable practical successes, leaving most of the deep scientific questions unanswered.

I am not claiming that progress is impossible, only that it is very difficult and requires integration across disciplines, and a long term strategy. It also depends on a very broad and deep educational system for potential high calibre researchers. That sort of education may not be forthcoming because of the likely cost, and the ignorance of planners at national levels.

Despite the enormous practical importance of developments in AI, there have always been some AI researchers who are more interested, in the potential of AI as **science** and **philosophy** than in the potential of AI systems to do useful things. In particular AI (along with computer science and computer systems engineering -- hardware and software) has begun to advance scientific and philosophical insights by providing new forms of explanation for aspects of natural intelligence and new answers to ancient philosophical questions about the nature of minds, their activities, and their products. In particular the deepest aim of science (not always acknowledged as such) is to discover what sorts of things are *possible*, and what makes, or could make, them possible. Explaining the possibility of some of the most complex organisms, including their information-processing competences, requires not only knowledge of biology, physics and chemistry, but also deep, highly advanced, engineering competences combined with philosophical knowledge and expertise. Our educational system mostly fails to achieve this.

Major scientific theories all contribute to the study of what is possible and how it is possible, including the ancient atomic theory, Newton's mechanics, chemistry, Darwin's theory of natural selection, quantum physics (see [Schrödinger\(1944\)](#)), computer science and AI (as explained in [Sloman\(1978b\)](#), [Chapter 2](#)). The Turing-inspired Meta-Morphogenesis project mentioned in [\[2\]](#) has been a part of this since 2012. AI, including future forms of AI, must be an essential part of any deep study of "the space of possible minds" [Sloman\(1984\)](#).

My own interest, for nearly half a century has centred mainly on the potential of AI to answer scientific and philosophical questions, e.g. about what minds and mental states and processes are, and how they work, including how they evolved, how they develop, how they are implemented in known physical/chemical mechanisms, and how we can use the new understanding to improve ways of helping people, for example in education, and therapy. A particular scientific sub-task is to explain how biological evolution is able to produce so many different forms of (more or less intelligent) information processing, in humans and non-human animals, and in humans at different stages of development, in different physical and cultural contexts, and in different cooperating subsystems within complex individuals (e.g. information processing subsystems involved in language development, visual perception, motivational processes, and mathematical discovery). Explaining all this requires advances in understanding of many varieties of information processing,

Some important clues may come from earlier evolutionary stages, including: microbe minds, insect minds, and evolutionary precursors of the more complex minds we hope to understand. Studying evolutionary transitions can provide hints as to changes in information processing mechanisms. This study is the Meta-Morphogenesis project mentioned in [Note \[2\]](#).

Unfortunately much "standard" scientific research that seeks experimental or naturally occurring regularities fails to identify what really needs to be explained: e.g. because most of what goes on in animal information processing is far richer than observable and repeatable input-output relationships -- e.g. your mental processes as you read this. No amount of actual laboratory testing can exhaust the responses you could possibly give to possible questions about what you are reading here, and there is no reason to assume that all humans, even from the same social group, or even the same research department, will give the same answers, and not only because of their different histories.

A standard, implicit, response is to regard all that diversity as irrelevant to a science of mind. One consequence of that attitude is research using experiments, e.g. in developmental psychology, designed to constrain subjects artificially to support repeatability, which can conceal their true potential, and a shortage of long term studies of individuals, which would have to accommodate enormous variability in developmental trajectories.

There are exceptions, e.g. Piaget's pioneering work, e.g. on [Possibility and Necessity](#), published posthumously. But he lacked adequate theories of information processing mechanisms (as he admitted at a workshop, shortly before he died). Piaget's earlier work inspired the educational proposals in [Sauvy and Sauvy \(1974\)](#). They could also provide useful tests for future, more human-like, robots.

Possibilities, not just laws, need to be explained

Overcoming the limitations of "standard" empirical research on how minds work requires setting explanatory goals at the level of *generative powers* rather than *observed regularities*, as Chomsky and others pointed out long ago [Chomsky\(1965\)](#). (For historical detail see [Boden\(2006\)](#). Compare the claim above that deep science is more concerned with discovery and explanation of *possibilities* than *laws* [Sloman\(1978b\) Chapter 2.](#))

Modelling observed regularities can often be achieved without accurate modelling of the mechanisms that happened on that occasion to produce those regularities, even in the physical sciences, e.g. the apparent successes of the Ptolemaic theory of planetary motion, and many other well supported then later abandoned regularities in physics -- including Newtonian dynamics.

The problems of reliance only on observed and repeatable regularities are far worse in the science of mind. Overcoming those problems requires application of deep multi-disciplinary knowledge and expertise, including the kind of expertise involved in designing testing and debugging complex virtual machines interacting with complex environments. (This helps to debunk the myth that AI is dependent on Turing machines: TMs are defined to run disconnected from any environment, rendering them useless for working AI systems, despite their great theoretical importance for computer science [Sloman\(2002\)](#). That paper also explains why using a Turing machine to implement multiple concurrently active interacting virtual machines is inherently unreliable.) Some preliminary suggestions regarding a "Super Turing membrane machine" are under development in Sloman [2017b](#) and [2017c](#), related to ideas in Sloman [\(2008\)](#) and [McClelland\(2017\)](#)'s work on affordances for mental action.

Insights can often be gained by studying naturally occurring, but relatively rare phenomena. For example, attempts to teach deaf children in Nicaragua to use sign language demonstrated that children do not merely [learn](#) pre-existing languages: they can also [create](#) new languages cooperatively, though this is cloaked by the fact that they are usually in a minority, so that collaborative construction looks like learning [Senghas \(2005\)](#). Close observation of competences of pre-verbal children and other intelligent species also provides evidence that richly structured languages must have evolved for internal use (e.g. for perceptual contents, intentions, questions, and planned complex actions) before external languages developed [Sloman\(1978a\)](#), [Sloman\(2015a\)](#). The same point is expressed in different words in [Mumford\(2016\)](#).

An example: explaining human/animal mathematical competences

A particular generative aspect of human intelligence that has been of interest to philosophers for centuries, discussed in [Kant\(1781\)](#), is the ability to make mathematical discoveries, including the amazing discoveries in geometry presented in Euclid's *Elements* over two thousand years ago that are still in use world-wide every day by scientists, engineers and mathematicians (though unfortunately now often taught only as facts to be memorised rather than rediscovered by learners).

I suspect that Kant understood that those abilities were deeply connected with practical abilities in non-mathematicians such as weaver birds, squirrels, elephants, and pre-verbal toddlers (my examples, not his), as illustrated in the video presentation in [Sloman 2017b](#). Young children don't have to be taught topology in order to understand that something is wrong when a stage magician appears to link and unlink a pair of solid metal rings. Some of the details are elaborated in [Sloman 2017c](#) (still under development at the time of writing this).

Despite the popular assumption that computers are particularly good at doing mathematics, because they can calculate so fast, run mathematical simulations, and even discover new theorems and new proofs of old theorems using AI theorem-proving packages, they still cannot replicate the ancient geometric and topological discoveries, or related discoveries of aspects of geometry and topology made unwittingly by human toddlers (illustrated in a video referenced in [Note\[4\]](#)) and related achievements of other species, e.g. birds that weave nests from twigs or leaves, and squirrels that defeat "squirrel-proof" bird feeders. (Search online for videos.)


These limits of computers are of far deeper significance for the science of minds than debates about whether computer-based systems can understand proofs of incompleteness theorems by Gödel and others, e.g. [Penrose \(1994\)](#). (Penrose recognized the importance of ancient geometric competences, but gave no plausible reasons to think they [cannot](#) be replicated in AI systems,

although they have not been replicated so far.)

There are impressive AI geometry theorem provers, but they [start](#) from logical formalisations of Euclid's axioms and postulates, and derive theorems from them using methods of modern logic, algebra, and arithmetic (e.g. in detecting false conjectures to prune search paths). Those methods are at most a few hundred years old, and some much less than that. They were not known to or used by great ancient mathematicians, such as Archimedes, Euclid, Pythagoras and Zeno, or children of my generation learning to prove statements in Euclidean geometry.

A major unsolved problem for AI is to understand and replicate the relevant reasoning powers. In particular, the postulates and axioms in Euclid's *Elements*, e.g. concerning congruency were stated without proof, but were not *arbitrary assumptions* adopted as starting points to define a mathematical domain, as in modern axiomatic systems.

Rather, Euclid's axioms and postulates were *major discoveries*, and various mathematicians and philosophers have investigated ways of deriving them from supposedly more primitive assumptions, e.g. deriving notions like point and line from more primitive spatial/topological notions (as demonstrated in [Scott\(2014\)](#). Here is a simpler example, from [Sloman\(2017b\)](#), elaborated in [Sloman\(2017c\)](#).



If you start with an arbitrary planar triangle, like the blue one then move its top vertex further from the opposite side, along a line through the opposite side, e.g. producing the red triangle, and then continuing, what happens to the size of the angle at the top as it moves: how do you know? What enables you to know that it is impossible for the angle to get larger?

Discovery of a more complex related problem is left as an exercise for readers -- alternatively read (and criticise!) this: [Sloman\(2017c\)](#).

Euclid's *starting points* require mathematical discovery mechanisms that seem to have gone unnoticed, and are not easily implementable in current AI systems without using something like a Cartesian-coordinate-based arithmetic model for geometry, which was not used by the ancient mathematicians making discoveries thousands of years before Descartes.

Moreover, for reasons given by Kant, they cannot be *empirical* discovery methods based on finding a regularity in many trial cases, since that cannot prove *impossibility*: mathematics is concerned with necessary truths and impossibilities not empirical generalisations. This does not imply *infallibility*, as shown by Lakatos (1976). Any practising mathematician knows that mathematicians can make mistakes. I did when reasoning about the stretched triangle problem above, which is what led to the exploration reported in another paper Sloman(2017c).

Is it possible to add the ancient mathematical discovery mechanisms to AI using current computing technology, or are new kinds of computers required, e.g. perhaps chemical computers replicating ill-understood brain mechanisms? (I suspect Turing was thinking about such mechanisms around the time he died (suggested by reading Turing(1952)). There is evidence that Kenneth Craik (1943), another who died tragically young, was also thinking about such matters, perhaps ahead of Turing.

Does any current neuroscientist understand how biological brain mechanisms can represent and reason about perfectly straight, perfectly thin lines, and their intersections? And reason about effects of moving them in a plane or on other surfaces?

Later work will need to dig deeper into similarities and differences between the forms of logical/mathematical reasoning that computers can or cannot cope with, e.g. because the former use manipulation of discrete structures or discrete search spaces, and the latter require new forms of computation, e.g. the structures and processes used in ancient proofs of geometrical and topological theorems.

Compare the procedures for deriving Euclid's ontology from geometry without points presented in a recorded lecture by Dana Scott (2014), using diagrammatic reasoning rather than logical and arithmetic reasoning.

The required new mechanisms are not restricted to esoteric activities of mathematical researchers: many non-mathematicians, including young children, find it obvious that two linked rings made of rigid impenetrable material cannot become unlinked without producing a gap in one of the rings.

How is such impossibility represented in animal brains? How are such impossibilities *derived* from perceived structural relationships? Young children don't have to study topology to realise that something is wrong when a stage magician appears to link and unlink solid rings. What mechanisms do their brains use? Or the brains of squirrels mentioned above?

Additional examples are presented here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html> There are many more examples, including aspects of everyday reasoning about clothing, furniture, effects of various kinds of motion, etc. and selection between possible actions (affordances) by using partial orderings in space during visual feedback rather than numerical measures of spatial relationships or using the kind of statistical/probabilistic reasoning that now (unfortunately) dominates AI work in vision and robotics. An alternative approach using semi-metrical reasoning, including topological relations and partial orderings, is suggested in

Sloman (2007-14).

Current computers can produce realistic simulations of **particular** spatial processes but that's very different from understanding generic constraints on **classes of processes**, like the fact mentioned above that if a vertex of a triangle moves away from the opposite side, whose length is fixed, then if the line passes between the other two angles, the angle at the moving vertex must always decrease in size during the motion.

No amount of repetition of such processes using a drawing package on a computer will enable the computer to understand **why** the angle gets smaller, or think of asking whether the monotonicity depends both on the choice of the line of motion of the vertex and the starting point, as discussed in Sloman(2017c).

Such geometric reasoning about partial orderings is very different from understanding why an expression in boolean logic is unsatisfiable or why a logical formula is not derivable from a given set of axioms, both of which can be achieved (in some cases) by current AI systems. It is also different from reasoning about the truths of arithmetical formulae corresponding to the geometrical structures and processes via use of Cartesian coordinates for points, lines and circles. (Objections by Searle and others that computers cannot understand anything have been adequately refuted elsewhere.)

Can we give the required sort of consciousness of geometrical necessity to future robots? The lack of any discussion of mathematical consciousness, e.g. "topological impossibility qualia" in most contemporary theories of consciousness, seems to me to suggest that those theories are at best incomplete, and probably deeply mistaken, at least as regards spatial consciousness.

The tendency for philosophers of mind to ignore mathematical discovery is particularly puzzling given the importance Kant attributed to the problem as long ago as 1781. (And long before him Socrates and Plato?)

Perhaps this omission is a result of a mistaken belief that Kant was proved wrong when empirical support was found for Einstein's claim that physical space is non-Euclidean. Had he known about non-Euclidean geometries Kant could have given as an example of non-empirical discovery of non-analytic mathematical truths the discovery that a subset of Euclidean geometry can be extended in different ways, yielding Euclidean and non-Euclidean geometries. Kant had no need to claim that human mathematicians are infallible, and as far as I know, never did claim that. His deep insights were qualified, not refuted, by Lakatos 1976.

Additional examples of types of mathematical and non-mathematical reasoning that need to be explained and modelled are presented in Sloman(2015-impossible). Some discoveries of that kind seem to be made (and used) by pre-verbal human toddlers, illustrated in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>

Whether AI can be extended in the foreseeable future to accommodate the ancient mathematical competences using current computers depends on whether we can implement the required virtual machinery in digital computers or whether, like brains, future human-like computers will have to make significant use of chemical information processing, using molecules rather than neurons as processing units Grant(2010), Trettenbrein(2016)

As long ago as 1944 ([Schrödinger](#)) pointed out the importance for life of the fact that quantum physics explains how chemistry can support both discrete processes (structural changes in chemical bonds) and continuous changes (folding, twisting, etc.) The possibility that biological information processing is implemented not at the neural level but at the molecular level was also considered by John von Neumann in his 1958 book *The computer and the brain*, written while he was dying. If true this implies that current calculations regarding how soon digital computers will replicate brain functionality are out by orders of magnitude (e.g. centuries rather than decades). See also [Newport\(2015\)](#).

AI researchers who have not studied Kant's views on the nature of mathematical knowledge as non-analytic (synthetic, i.e. not derivable using only definitions and pure logic), non-contingent (concerned with what's possible, necessarily the case, or impossible) may find it hard to understand what's missing from AI. In particular, I have found that some believe that eventually deep learning mechanisms will suffice.

But mechanisms using only statistical information and probabilistic reasoning are constitutionally incapable of learning about necessary truths and falsehoods, as Kant noticed, long ago, when he objected to Hume's claim that there are only two kinds of knowledge: empirical knowledge and analytic knowledge (definitional relations between ideas, and their logical consequences).

Hume's view of [causation](#) as being of the first sort (concerned with observed regularities) is contradicted by mathematical examples including the triangle deformation example above: motion of a vertex of a triangle away from the opposite side [causes](#) the angle to decrease, just as adding three apples to a collection of five apples causes the number in the collection to increase to eight. (Examples of Humean and Kantian causal reasoning in humans and possibly other animals are presented in [Sloman and Chappell \(2007b\)](#).)

Can AI lead to robots with these ancient mathematical reasoning abilities?

I'll indicate possible lines of enquiry to discover what's missing from current AI, partly inspired by asking where Turing was heading in his 1952 paper, and partly based on a new theory regarding the variety of mechanisms and transitions in biological evolution, including the evolution of many new kinds of construction kit [Sloman\(2017\)](#), many of which introduced new kinds of information processing mechanism, and new ideas about epigenetic processes that could produce young potential mathematicians. (Some of the ideas were developed in collaboration with biologist Jackie Chappell [Chappell & Sloman \(2007a\)](#) (i.e. the "Meta-configured genome") a decade ago, related to extended versions of Karmiloff-Smith's theories of "Representational Redescription" in (1992) .

One consequence of these investigations is rejection of the popular "Possible worlds semantics" as an analysis of (alethic) modal operators: "impossible", "possible", "contingent", and "necessary", in favour of (Kant-inspired) semantics related to variations in configurations of fragments of *this* world, as illustrated in the stretched triangle example. Other implications for AI as Science, AI as engineering and AI as philosophy, will be discussed. There may or may not be time to present detailed ideas about a super-Turing membrane computer, hinted at above.

NOTES

1. An invited talk to a workshop at IJCAI 2017 (August 2017), closely related to this topic, was presented in a pre-recorded 42 minute video, available online, along with extended notes, here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html> Since the event those notes have been substantially extended.
 2. This is part of an ambitious Turing-inspired project investigating the evolution of biological information processing mechanisms and capabilities, the Meta-Morphogenesis project <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>, which includes a new theory of evolved construction-kits [Sloman\(2017\)](#).
 3. E.g. in the epilogue to my 1978 book, *The Computer Revolution in Philosophy*, here <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#epilogue>
 4. 3D Toddler topology is illustrated in this 4.5min video (with commentary): [WEBM\[about 4 minutes\]](#)
-

REFERENCES (A small sample of relevant literature)

- M. A. Boden, 1978, *Artificial Intelligence and Natural Man*, Harvester Press, Hassocks, Second edition 1986. MIT Press,
- M. A. Boden, 2006, *Mind As Machine: A history of Cognitive Science* (Vols 1--2) OUP
- Jackie Chappell and Aaron Sloman, 2007, Natural and artificial meta-configured altricial information-processing systems, in *International Journal of Unconventional Computing*, 3, 3, pp. 211--239, <http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717>
- N. Chomsky, 1965, *Aspects of the theory of syntax*, MIT Press, Cambridge, MA.
- Max Clowes (1967) Perception, picture processing and computers, in *Machine Intelligence Vol 1* pp 181--197, Eds. N L Collins and Donald Michie, Oliver & Boyd, <http://www.cs.bham.ac.uk/research/projects/cogaff/MI1-2-Ch.12-Clowes.pdf>
- Shang-Ching Chou, Xiao-Shan Gao and Jing-Zhong Zhang, 1994, *Machine Proofs In Geometry: Automated Production of Readable Proofs for Geometry Theorems*, World Scientific, Singapore, <http://www.mmrc.iss.ac.cn/~xgao/paper/book-area.pdf>
- D. C. Dennett, 1978 *Brainstorms: Philosophical Essays on Mind and Psychology*, MIT Press, Cambridge, MA,
- D. C. Dennett, 1984, *Elbow Room: the varieties of free will worth wanting*, Oxford: The Clarendon Press,
- D.C. Dennett, 1996
Kinds of minds: towards an understanding of consciousness,
Weidenfeld and Nicholson, London, 1996,

- Shannon Densmore and Daniel Dennett, 1999 The Virtues of Virtual Machines, pp. 747--761, *Philosophy and Phenomenological Research*, 59, 3, Sep, 1999, <http://www.jstor.org/stable/i345616>
- Kenneth Craik, 1943, *The Nature of Explanation*, Cambridge University Press, London, New York,
- J. J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, MA, 1979.
- H. Gelernter, 1964, Realization of a geometry-theorem proving machine, pp. 134-152, in *Computers and Thought*, Eds. E. Feigenbaum and J. Feldman, McGraw-Hill, New York,
- Seth Grant, 2010, Computing behaviour in complex synapses - synapse proteome complexity and the evolution of behaviour and disease, *Biochemical Society Magazine*, Vol 32 NO 2 April 2010 <http://www.biochemist.org/bio/default.htm?VOL=32&ISSUE=2>
- Immanuel Kant, 1781, *Critique of Pure Reason*, Translated (1929) by Norman Kemp Smith, London, Macmillan,
- I. Lakatos, 1976, *Proofs and Refutations*, Cambridge University Press, Cambridge, UK,
- John McCarthy and Patrick J. Hayes, 1969, "Some philosophical problems from the standpoint of AI", *Machine Intelligence 4*, Eds. B. Meltzer and D. Michie, pp. 463--502, Edinburgh University Press, <http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html>
- J. McCarthy, "The well-designed child", *Artificial Intelligence*, 172(18), 2003-2014, (2008). (Originally written in 1996) <http://www-formal.stanford.edu/jmc/child.html>
- Tom McClelland, (2017) AI and affordances for mental action, in Computing and Philosophy Symposium, *Proceedings of the AISB Annual Convention 2017* pp. 372-379. April 2017. <http://wrap.warwick.ac.uk/87246>
- A. Karmiloff-Smith, *Beyond Modularity: A Developmental Perspective on Cognitive Science*, MIT Press, Cambridge, MA, 1992.
- Colette Maloney (2003) Cognitive Systems scope and objectives Information Day Luxembourg, 20.06.2003 http://cordis.europa.eu/fp7/ict/robotics/docs/maloney-jun2003_en.pdf
An archive of reports and presentations relating to the Cognitive Systems project, in reverse chronological order is here http://cordis.europa.eu/fp7/ict/robotics/past-calls_en.html
- COMPARE: UKCRC Grand challenges 5 and 7 (2002--)
<http://www.cs.stir.ac.uk/gc5/>
<https://www.cs.york.ac.uk/nature/gc7/>
- G.A. Miller, E. Galanter and K.H. Pribram, 1960, *Plans and the Structure of Behaviour* Holt, New York,
- Ed. M. L. Minsky, 1968, Matter Mind and Models, in *Semantic Information Processing*, MIT Press, Cambridge, MA,

David Mumford, Grammar isn't merely part of language, Oct, 2016, (Online Blog)
<http://www.dam.brown.edu/people/mumford/blog/2016/grammar.html>

Tuck Newport, 2015 *Brains and Computers: Amino Acids versus Transistors*,
<https://www.amazon.com/dp/B00OQFN6LA>

Roger Penrose (1994), *Shadows of the mind: A Search for the Missing Science of Consciousness*.
OUP, Oxford

Jean Piaget (1981.1983), *Possibility and Necessity*

Vol 1. The role of possibility in cognitive development (1981)

Vol 2. The role of necessity in cognitive development (1983)

Tr. by Helga Feider from French in 1987

(Like Kant, Piaget had deep observations but lacked an understanding of information processing mechanisms, required for explanatory theories.)

Jean Sauvy and Simonne Sauvy, with an introduction by Bill Brookes *The Child's Discovery of Space: From hopscotch to mazes -- an introduction to intuitive topology*,
Penguin Education, Harmondsworth, 1974. Translated from the French by Pam Wells.

Dana Scott, 2014, Geometry without points. (Video lecture, 23 June 2014, University of Edinburgh)
<https://www.youtube.com/watch?v=sDGnE8eja5o>

Schmidhuber, J., 2014, Deep Learning in Neural Networks: An Overview, Technical Report
IDSIA-03-14, <http://arxiv.org/abs/1404.7828>

Erwin Schrödinger, *What is life?*, CUP, Cambridge, 1944.

Commented extracts available here:

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

A. Senghas (2005). Language Emergence: Clues from a New Bedouin Sign Language. *Current Biology*, 15 (12), R463-R465, Elsevier <http://dx.doi.org/10.1016/j.cub.2005.06.018>

See also this compelling BBC video: The Birth of New Sign Language in Nicaragua

<http://www.youtube.com/watch?v=pjtiolFuNf8>

H. A. Simon, 1967, Motivational and emotional controls of cognition, reprinted in *Models of Thought*, Ed. H. A. Simon, Yale University Press, pp. 29--38, Newhaven, CT,

A. Sloman, 1971, "Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence", in *Proc 2nd IJCAI*, pp. 209--226, London. William Kaufmann. Reprinted in *Artificial Intelligence*, vol 2, 3-4, pp 209-225, 1971.

<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#1971-02>

A. Sloman, (1978a), What About Their Internal Languages? Commentary on three articles by Premack, D., Woodruff, G., by Griffin, D.R., and by Savage-Rumbaugh, E.S., Rumbaugh, D.R., Boysen, S. in *Behavioral and Brain Sciences Journal* 1978, 1 (4),

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#713>

A. Sloman (1978b), *The Computer Revolution in Philosophy, Philosophy, Science and Models of Mind*, Harvester Press.

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/>

A. Sloman, 1984, The structure of the space of possible minds, in *The Mind and the Machine: philosophical aspects of Artificial Intelligence*, Ed. S. Torrance, Ellis Horwood, Chichester,

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#49a>

A. Sloman, 1986 Did Searle attack strong strong or weak strong AI, Ed. A.G. Cohn and J.R. Thomas, *Artificial Intelligence and Its Applications*, John Wiley and Sons,

<http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#70>

A. Sloman (2002), The irrelevance of Turing machines to AI, in *Computationalism: New Directions*, Ed. M. Scheutz, MIT Press, Cambridge, MA, pp. 87--127,

<http://www.cs.bham.ac.uk/research/cogaff/00-02.html#77>

Aaron Sloman (2007-2014)

Discussion Paper: Predicting Affordance Changes: Steps towards knowledge-based visual servoing. (Including videos).

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

Aaron Sloman, (2008), Architectural and Representational Requirements for Seeing Processes, Proto-affordances and Affordances, in *Logic and Probability for Scene Interpretation*, Eds. A.G. Cohn, D.C. Hogg, Ralf Moeller and Bernd Neumann, Dagstuhl Seminar Proceedings, No 08091, Schloss Dagstuhl Germany, 2008, <http://drops.dagstuhl.de/opus/volltexte/2008/1656/>

A. Sloman (2015a). What are the functions of vision? How did human language evolve? Online research presentation.

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>

A. Sloman 2017a, "Construction kits for evolving life (Including evolving minds and mathematical abilities.)" in *The Incomputable Journeys Beyond the Turing Barrier*

Eds: S. Barry Cooper and Mariya I. Soskova

<https://link.springer.com/book/10.1007/978-3-319-43669-2>

Unpublished extended (still growing) version:

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

A. Sloman 2017b Why can't (current) machines reason like Euclid or even human toddlers? (And many other intelligent animals), invited presentation for Workshop on Architectures for Generality and Autonomy <http://cadia.ru.is/workshops/aga2017/> at IJCAI 2017, Melbourne, August 2017.

Video recording and online notes available

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html>

A. Sloman 2017c Non-Monotonic angle size change as a vertex moves on a line. Online discussion paper, School of Computer Science, University of Birmingham. (Still under development.)

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

Aaron Sloman and Jackie Chappell (2007b)

Humean and Kantian causal reasoning in animals and machines.

Two ways of understanding causation: Humean and Kantian (linked invited presentations) at WONAC: International Workshop on Natural and Artificial Cognition, Pembroke College, Oxford June 2007. <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac>

A. Sloman and David Vernon. A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots, 2007. Contribution to euCognition wiki.

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

Max Tegmark, 2014, *Our mathematical universe, my quest for the ultimate nature of reality*, Knopf (USA) Allen Lane (UK), (ISBN 978-0307599803/978-1846144769)

Trettenbrein, Patrick C., 2016, The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift?, *Frontiers in Systems Neuroscience*, Vol 88,

<http://doi.org/10.3389/fnsys.2016.00088>

Trettenbrein, Patrick C., 2016, The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift?, *Frontiers in Systems Neuroscience*, Vol 88,

<http://doi.org/10.3389/fnsys.2016.00088>

A. M. Turing, 'Computing machinery and intelligence', *Mind*, 59, 433-460, (1950). (reprinted in E.A. Feigenbaum and J. Feldman (eds) *Computers and Thought* McGraw-Hill, New York, 1963, 11-35).

A. M. Turing, (1952) "The Chemical Basis Of Morphogenesis", *Phil. Trans. Royal Soc. London B* 237, 237, 37-72.

Note: A presentation of the main ideas for non-mathematicians can be found in Philip Ball, 2015, "Forging patterns and making waves from biology to geology: a commentary on Turing (1952) 'The chemical basis of morphogenesis'",

<http://dx.doi.org/10.1098/rstb.2014.0218>

Barbara Vetter (2011), Recent Work: Modality without Possible Worlds, *Analysis*, 71, 4, pp. 742--754, <https://doi.org/10.1093/analys/anr077>

C. H. Waddington, *The Strategy of the Genes. A Discussion of Some Aspects of Theoretical Biology*, George Allen & Unwin, 1957.

L. Wittgenstein (1956), *Remarks on the Foundations of Mathematics*, translated from German by G.E.M. Anscombe, edited by G.H. von Wright and Rush Rhees, first published in 1956, by Blackwell, Oxford. There are later editions. (1978: VII 33, p. 399)
