Huge, Unnoticed, Gaps Between Current AI and Natural Intelligence

Aaron Sloman
http://www.cs.bham.ac.uk/~axs*

October 12, 2018

Abstract

Despite AI's enormous *practical* successes, some researchers focus on its potential as science and philosophy: providing answers to ancient questions about what minds are, how they work, how multiple varieties of minds can be produced by biological evolution, including minds at different stages of evolution, and different stages of development in individual organisms. AI cannot yet replicate or faithfully model most of these, including ancient, but still widely used, mathematical discoveries described by Kant as non-empirical, non-logical and non-contingent. Automated geometric theorem provers start from externally provided logical axioms, whereas for ancient mathematicians the axioms in Euclid's *Elements* were major *discoveries*, not arbitrary starting points. Human toddlers and other animals spontaneously make similar but simpler topological and geometrical discoveries, and use them in forming intentions and planning or controlling actions. The ancient mathematical discoveries were not results of statistical/probabilistic learning, because, as noted by Kant, they provide non-empirical knowledge of possibilities, impossibilities and necessary connections. Can gaps between natural and artificial reasoning in topology and geometry be bridged if future AI systems use previously unknown forms of information processing machinery – perhaps "Super-Turing Multi-Membrane" machinery?

Keywords/phrases:
AI as science and philosophy; Can AI model ancient geometers? Can AI model human toddlers? Gaps and limitations of current AI; Super-Turing membrane machines; Replicating mathematical consciousness; Research needed.

1 The Meta-Morphogenesis Project

This paper opens a small window into a large project, begun over half a century ago, in my DPhil thesis (Sloman, 1962) defending Kant's claims (Kant, 1781) about the nature of

*Preprint for proceedings PT-AI-17, Ed. Vincent Müller, SAPERE 44, pp. 92-105, 2018. https://doi.org/10.1007/978-3-319-96448-5_11 Springer

mathematical discoveries: they are non-empirical, non-contingent, and they are *synthetic*, i.e. not based purely on logic plus definitions.

Around 1969 Max Clowes introduced me to AI. The project then grew into an attempt to use AI to explain many aspects of minds (Sloman, 1978 Revised 2018), including their abilities to make mathematical discoveries, especially the geometrical and topological discoveries made by ancient mathematicians.

A new strand began in 2011, inspired by Turing's work on morphogenesis (Turing, 1952), namely the Meta-Morphogenesis (M-M) project[1], investigating evolution of biological information processing mechanisms and capabilities, including an outline theory of evolved construction-kits.[2,3] That provides a framework for new attempts to identify ancient processes and mechanisms of mathematical discovery, especially precursors of mechanisms involved in topological and geometric discovery, illustrated by the work of Archimedes, Euclid, Zeno and many others. I conjecture that this will require discovery of some of the simpler intermediate cases in the evolution of the mechanisms involved. Early developmental stages may also give clues.

Studying spatial reasoning in other intelligent species, e.g. squirrels and crows, and preverbal human toddlers, may give clues regarding mechanisms used by ancient adult human mathematicians, including clues indicating their reasoning about possible and impossible spatial structures and processes in solving practical problems, where those processes use subsets of the mechanisms involved in ancient mathematical discoveries.

There's no evidence that ancient mathematicians and intelligent non-human animals use axiomatic, logical, forms of representation and reasoning based on Cartesian coordinates, such as Hilbert's axiomatization of Euclid (Hilbert, 1899), and geometry theorem provers, e.g. (Chou, Gao, & Zhang, 1994). My claim could be challenged by evidence showing that brains of some non-human species, and humans who have never encountered modern logic include genetically specified formalisms and mechanisms for doing what logic theorem provers do. (Merely showing that activity in certain brain regions is correlated with performing a task does not explain how brains perform that task – unlike specifying the algorithms and data-structures used by a robot to perform the task.)

Analysing examples of related, simpler, mathematical and proto-mathematical discoveries in humans and other animals[4], suggests that intelligent animals use types of information processing machinery that are not included in currently understood logical, algebraic, or statistical, reasoning mechanisms, including neural-nets. For example, no learning mechanism based on probabilistic inference can discover impossibilities or necessities, which are key features of mathematical discovery, as pointed out in (Kant, 1781).

Virtual machines running on digital computers closely coupled with the environment could be richer than a Turing machine, e.g. if the environment includes non-digital or truly

[1] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html

[2] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

[3] An invited video talk at IJCAI 2017, is available online, with extended notes:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html

[4] Pre-verbal toddler topology is illustrated in this 4.5min video:
http:www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/small-pencil-vid.webm

random phenomena). If the environment with which a digital computer interacts is not a discrete-state machine, the coupled system, including any virtual machinery used, cannot be modelled with perfect precision on a Turing machine, since no discrete machine can model perfectly a processes that runs through all the real numbers between 1 and 2, in order, whereas a continuously changing chemical structure might be able to.[5] (Moving only through the rationals in order would be *far* more complex. Why?)

2 Limited progress, despite spectacular successes

The practical uses of AI, and the rate at which they are now multiplying are 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[6], that truly superhuman intelligence will include a kind of wisdom that rejects the selfish, thoughtless, competitive, destructive, gullible, 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: 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 a new deep understanding of natural intelligence. For a survey see Margaret Boden's two-volume masterpiece (2006).

Recent spectacular engineering successes mask (current) limited scientific and philosophical progress in AI. Two results of this masking (at present) are a shortage of good researchers focusing on the long term issues, and a shortage of funds for long term scientific research. Most funded AI research at present aims at demonstrable practical successes, leaving some of the important scientific questions unanswered, and to some extent un-noticed!

I do not claim that progress is impossible, only that it is very difficult and requires deep integration across disciplines. It also depends on an educational system producing high calibre multi-disciplinary researchers.

Despite its enormous practical importance, some AI researchers, like Turing, are more interested in the potential of AI as *science* and *philosophy* than its *practical* applications. E.g. AI (along with computer science) has begun to advance science and philosophy 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, as explained in Chapter 2 of (Sloman, 1978 Revised 2018), the deepest aim of science (not always acknowledged as such) is to discover what sorts of things are *possible*,

[5] For further discussion of "virtual machine functionalism" see http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html

[6] 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

and what makes, or could make, them possible, not to discover regularities. In contrast, many science students are (unfortunately) taught to regard science as primarily concerned with finding, explaining and using observed correlations: a shallow view of science criticised vehemently in (Deutsch, 2011). Deep scientific theories all contribute to the study of what is possible and how it is possible, including ancient atomic theory, Newton's mechanics, chemistry, Darwin's theory of natural selection, quantum mechanics, e.g. (Schrödinger, 1944), computer science, AI and theoretical linguistics.

The Turing-inspired Meta-Morphogenesis project mentioned in Note 1 has addressed such issues 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), which may be far richer than anyone currently suspects.

3 AI as Science and Philosophy

For most people, AI is primarily an engineering activity, whereas my interest, since around 1969, inspired by Max Clowes, and AI founders such as Minsky e.g. (1963, 1968, 2006), McCarthy e.g. (1979, 2008), and Simon, e.g. (1967, 1969), is focused mainly on the potential of AI to trigger and eventually 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 can vary, with potential applications in education and therapy.

A long term goal is to explain how biological evolution is able to produce so many different forms of information-processing, in humans and non-human organisms, 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: internal languages[7], language development, visual perception, motivational processes, and mathematical discovery). Explaining all this requires major progress in understanding varieties of information processing. Clues may come from many evolutionary stages, including: microbe minds, insect minds, and other precursors of the most complex minds we hope to understand and model. This is the Meta-Morphogenesis project mentioned in Note 1.

Unfortunately much "standard" research, seeking experimental or naturally occurring regularities, fails to identify what needs to be explained, because most 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 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. Compare how different the outputs of great composers, or poets, or novelists are, even if they live in the same location. A standard response is to regard all that diversity as irrelevant to a science of mind. One consequence is narrowly focused research using experiments, e.g. in developmental psychology, designed to constrain subjects artificially to support repeatability. This can conceal their true potential, requiring long term studies of individuals, which would have to accommodate enormous

[7] http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111

variability in developmental trajectories.

There are exceptions, e.g. Piaget's pioneering work on children's understanding of Possibility and Necessity, published posthumously (1981,1983). But he lacked adequate theories of information processing mechanisms (as he admitted at a workshop I attended, shortly before he died). Piaget's earlier work inspired the proposals in (Sauvy & Sauvy, 1974). It could also suggest useful goals for future, more human-like, robots.

4 Aim for generative power *not* data summaries

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 (1965). (For historical detail see (Boden, 2006); Compare the claim that deep science is more concerned with discovery and explanation of *possibilities* than *laws*, in (Sloman, 1978 Revised 2018, Chap2).

Even in the physical sciences, modelling observed regularities can often be achieved without accurate modelling of the *mechanisms* that happened, on that occasion, to produce those regularities, 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.

Problems of reliance only on observed and repeatable regularities are far worse in the science of mind. Overcoming them requires application of deep multi-disciplinary knowledge and expertise, including 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). Preliminary ideas regarding a "Super Turing membrane machine" are in (Sloman, 2017b),[8] related to ideas about affordances in (Sloman, 2008) and McClelland's work on affordances for mental action, e.g. (2017). This requires substantial long term research.

Insights can often be gained by studying naturally occurring, but relatively rare phenomena, for example when 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).

4.1 Human/animal mathematical competences

A particular generative aspect of human intelligence that has been of interest to philosophers for centuries, and discussed by Kant (1781, 1783), is the ability to make mathematical discoveries, including the amazing discoveries in geometry presented in Euclid's *Elements* over

[8] For a detailed example see
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html
and
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html

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. Online documents exploring some of the details are referenced in Note 8 and the work on evolved construction-kits in Note 2.

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 the 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) (who recognizes the importance of ancient geometric competences, but gives no plausible reasons to think they *cannot* be replicated in AI systems, although they have not been replicated so far.)

4.2 AI theorem provers do something different

There are impressive AI geometry theorem provers, but they *start* from logical formalisations of Euclid's axioms and postulates, e.g. using Hilbert's (1899) version. They derive theorems using methods of modern logic, algebra, and arithmetic (e.g. pruning search paths by using numerical checks). Those methods are at most a few hundred years old, and some much newer. 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. How did their brains work?

A major unsolved problem for AI is to understand and replicate the relevant *ancient* reasoning powers.

In Fig. 1 (below) what happens to the size of the angle at A if A is moved further from BC along a straight line through the opposite side BC? Answering the question involves thinking about two continua (the continuum of positions of the top vertex, and the continuum of angle sizes) and their relations. Many people with no mathematical training can do this easily, in my experience. What are their brains doing? How do brains represent impossibility or necessity? If the line of motion of A intersects the base *outside* the triangle the situation is more complicated, and Apollonius' construction becomes relevant, as Diana Sofronieva pointed out to me.

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

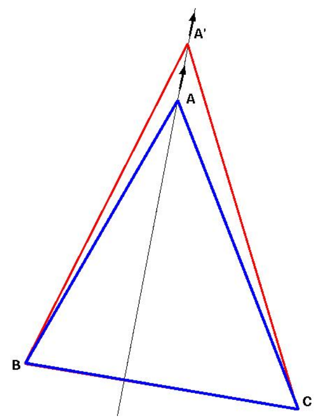

Figure 1:

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 by Dana Scott (2014). A simpler example, from (Sloman, 2017b), referenced in Note 8 is in Fig. 1.

If you start with an arbitrary planar triangle, like the blue one in Fig. 1, then continuously move one 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? Investigation of how the problem changes if the line of motion changes is left as an exercise for the reader (see Note 8).

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 only on finding regularities in many trial cases, since that cannot prove *necessity* or *impossibility*: mathematics is concerned with *necessary* truths and *impossibilities* not empirical generalisations. This feature is is ignored by much psychological research on mathematical competences and cannot be explained by statistics-based neural theories of mathematical reasoning. This does not imply *infallibility*, as shown by (Lakatos, 1976). Any practising mathematician knows that mathematicians can make mistakes. I did at first when reasoning about the stretched triangle problem above, which is what led to the exploration reported in (Sloman, 2017b).

5 Robots with ancient mathematical competences?

Can current computing technology support ancient mathematical discovery mechanisms, or are new kinds of computers required, e.g. perhaps chemical computers replicating illunderstood brain mechanisms? (I suspect Turing was thinking about such mechanisms around the time he died (suggested by reading (Turing, 1952)). There is evidence in (Craik, 1943) that Kenneth Craik, another who died tragically young, was also thinking about such matters, perhaps inspiring Turing posthumously? Does anything in current neuroscience explain how biological brain mechanisms represent and reason about *perfectly straight, perfectly thin* lines, and their intersections? Or reason about *impossibilities*, and *necessary* consequences of certain kinds of motion?

Future work needs to dig deeper into differences between the forms of logical/mathematical reasoning that computers can and 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 (Scott, 2014). The presentation clearly uses a great deal of spatial/diagrammatic reasoning rather than purely logical and algebraic reasoning.)

The required new mechanisms are not restricted to esoteric activities of mathematicians: e.g. 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.

6 Representing impossibility and necessity

What brain mechanisms can represent impossibility? How can impossibilities be*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?[9] 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 the kinds of statistical/probabilistic reasoning that now (unfortunately) dominate AI work in vision and robotics. An alternative approach uses semi-metrical reasoning, including topological structures and partial orderings, was suggested in (Sloman, 2007). I have not been able to persuade any AI/Robotics researchers, however, possibly because using that approach would require massive changes to Robot vision and reasoning mechanisms. How can such mechanisms be implemented in brains?

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 regularity linking two dimensions of continuous variation mentioned in Fig. 1.

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 to think of asking whether the monotonicity depends both on the choice of the line of motion of the vertex and the starting point. See Note 8 and (Sloman, 2017b). I did not notice this until Auke Booij

[9] Many additional examples are presented in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html.

pointed it out to me.

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, but only after the problem is rephrased in terms of possible sequences of logical formulae in a proof system, or possible solutions to numerical equations, using something like Hilbert's logic-based formulation of Euclidean geometry. Ancient geometric reasoning was very different from reasoning about arithmetical formulae by using Cartesian coordinates. (Claims by John Searle and others that computers are purely syntactic engines, with no semantic competences, have been adequately refuted elsewhere.)

7 Gaps in theories of consciousness

7.1 What is mathematical consciousness?

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 all contemporary theories of consciousness that I have encountered, 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 widely held, but mistaken, belief that Kant was proved wrong when empirical support was found for Einstein's claim that physical space is non-Euclidean. Had Kant known about non-Euclidean geometries, he could have given as his 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 different geometries with different properties. 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 (Lakatos, 1976). This was also discussed in my 1962 thesis (Sloman, 1962). "Proto-mathematical" discoveries of various kinds are also made, and put to practical uses, by pre-verbal human toddlers.[10]

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, perhaps using molecules rather than neurons as processing units, as discussed by Grant (Grant, 2010), Trettenbrein (Trettenbrein, 2016), Gallistel (Gallistel & Matzel, 2012), Newport (Newport, 2015) (citing von Neumann) and others.

As long ago as 1944 Schrödinger (Schrödinger, 1944) pointed out the importance for

[10]Several examples of various kinds are presented in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html

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 many orders of magnitude (e.g. many centuries rather than decades). See also (Newport, 2015).

7.2 Probabilistic reasoning *vs* impossibility/necessity

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 other animals were presented (in collaboration with Jackie Chappell) in (Chappell & Sloman, 2007b).

7.3 Can we give robots geometric reasoning abilities?

Possible lines of enquiry about what's missing from current AI are suggested by (Turing, 1952), leading to a new theory regarding the variety of mechanisms and transitions in biological evolution, including evolution of new kinds of construction kit (Sloman, 2017a).[11] Evolution repeatedly produced new biological construction kits for new kinds of information processing mechanism. This may explain the evolution of epigenetic processes that produce young potential mathematicians. Ideas about "meta-configured competences" are being developed in collaboration with biologist Jackie Chappell (Chappell & Sloman, 2007a),[12] extending Karmiloff-Smith's theories of "Representational Redescription" (1992), and hypotheses about non-linear, structured, extendable, *internal* languages required for percepts,

[11] The work on construction kits is still being extended in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

[12] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html

intentions, plans, usable generalisations, and reasoning, long before *external* languages were used for communication (Sloman, 2015).

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, and many other examples of geometrical and topological reasoning.

8 Conclusion, and further work

This paper opens a small window into a large, complex, still growing project. (See Note 1.) There are many implications for AI as Science, AI as engineering and AI as philosophy, and also deep implications for psychology and neuroscience, insofar as they have not yet addressed the problem of how minds or brains are able to make discoveries concerning necessary truths and impossibilities that are not merely logical truths or falsehoods. There are also hard biological problems to be solved, concerning evolutionary histories of the features of human brains and minds that have these amazing capabilities. Perhaps only after these non-AI questions have been answered will AI engineers be able to design artificial minds with ancient mathematical capabilities of Archimedes and others. Not all psychologists and neuroscientists notice that the task of explaining mathematical cognition is not merely the task of explaining *numerical* competences, on which they tend to focus, while ignoring the richness of numerical competences such as the central role of transitivity of one-one correspondence. (Piaget was an exception.)

On-going work investigates requirements for a Super-Turing membrane computer,[13] able to acquire and use information about spatial structures and relationships in performing practical tasks, for instance understanding how available information and affordances *necessarily* change as viewpoints change, or objects rotate or move – because visual information normally travels in straight lines. If these ideas can be used in future designs, we may be able to produce robots that replicate the discoveries made by great ancient mathematicians as well as the deep but unnoticed spatial reasoning abilities developed by pre-verbal humans and many other intelligent species.

This should help to stifle distracting and impoverished theories of embodied cognition, mistakenly giving the impression that there is no requirement for deep and complex internal information-processing engines produced by biological evolution, but not yet replicated in AI systems. And if, as I suspect (and perhaps Turing suspected), these mechanisms are implemented in sub-synaptic chemical mechanisms, then since there are many orders of magnitude more molecules than neurones, this suggests that hopes or fears about computers soon reaching or overtaking human intelligence are time-wasting distractions from the hard task of trying to understand and model human intelligence, or more generally animal intelligence.

The Meta-Morphogenesis web site is expected to continue growing.[14] But there are

[13] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-geom.html

[14] See Note 1, and also

many unsolved problems, including problems about mechanisms underlying ancient forms of mathematical consciousness.

References

Boden, M. A. (2006). *Mind As Machine: A history of Cognitive Science (Vols 1–2)*. Oxford: Oxford University Press.

Chappell, J., & Sloman, A. (2007a). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, *3*(3), 211–239. Retrieved from http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717

Chappell, J., & Sloman, A. (2007b, June 25-26). *Two ways of understanding causation: Humean and Kantian.* Pembroke College, Oxford. Retrieved from http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac

Chomsky, N. (1965). *Aspects of the theory of syntax*. Cambridge, MA: MIT.

Chou, S.-C., Gao, X.-S., & Zhang, J.-Z. (1994). *Machine Proofs In Geometry*. Singapore: World Scientific.

Craik, K. (1943). *The nature of explanation*. London, New York: CUP.

Deutsch, D. (2011). *The Beginning of Infinity: Explanations That Transform the World*. Allen Lane & Penguin Books.

Gallistel, C., & Matzel, L. (2012). The neuroscience of learning: beyond the Hebbian synapse. *Annual Revue of Psychology*, *64*, 169–200.

Grant, S. G. (2010, April). Computing behaviour in complex synapses. *Biochemist*, *32*, 6–9.

Hilbert, D. (1899). *The Foundations of Geometry*. Salt Lake City: Project Gutenberg. (Tr. 1902 by E.J. Townsend, from 1899 German edition)

Kant, I. (1781). *Critique of pure reason*. London: Macmillan. (Translated (1929) by Norman Kemp Smith)

Kant, I. (1783). *Prolegomena to Any Future Metaphysics That Will Be Able to Present Itself as a Science*. unknown. (Several translations online.)

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.

Lakatos, I. (1976). *Proofs and Refutations*. Cambridge, UK: CUP.

McCarthy, J. (1979). Ascribing mental qualities to machines. In M. Ringle (Ed.), *Philosophical perspectives in artificial intelligence* (pp. 161–195). Atlantic Highlands, NJ: Humanities Press.

McCarthy, J. (2008). The well-designed child. *Artificial Intelligence*, *172*(18),

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

2003–2014. Retrieved from http://www-formal.stanford.edu/jmc/child.html
McClelland, T. (2017, Apr 19–21). AI and affordances for mental action. In *Computing and Philosophy Symposium, AISB Convention 2017* (p. 372-37). Bath: AISB. Retrieved from http://wrap.warwick.ac.uk/87246
Minsky, M. L. (1963). Steps toward artificial intelligence. In E. Feigenbaum & J. Feldman (Eds.), *Computers and thought* (pp. 406–450). New York: McGraw-Hill.
Minsky, M. L. (1968). Matter Mind and Models. In M. L. Minsky (Ed.), *Semantic Information Processing*. Cambridge, MA: MIT Press.
Minsky, M. L. (2006). *The Emotion Machine*. New York: Pantheon.
Newport, T. (2015). *Brains and Computers: Amino Acids versus Transistors*. Kindle. Retrieved from https://www.amazon.com/dp/B00OQFN6LA
Penrose, R. (1994). *Shadows of the mind: A Search for the Missing Science of Consciousness*. Oxford: OUP.
Piaget, J. (1981,1983). *Possibility and Necessity: Vol 1. The role of possibility in cognitive development, Vol 2. The role of necessity in cognitive development*. Minneapolis: U. of Minnesota Press. (Tr. by Helga Feider from French in 1987)
Sauvy, J., & Sauvy, S. (1974). *The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology*. Harmondsworth: Penguin Education. (Translated from the French by Pam Wells)
Schrödinger, E. (1944). *What is life?* Cambridge: CUP.
Scott, D. S. (2014). *Geometry without points.* Retrieved from https://www.logic.at/latd2014/2014%20Vienna%20Scott.pdf
Senghas, A. (2005). Language Emergence: Clues from a New Bedouin Sign Language. *Current Biology*, *15*(12), R463–R465.
Simon, H. A. (1967). Motivational and emotional controls of cognition. In H. A. Simon (Ed.), *reprinted in models of thought* (pp. 29–38). Newhaven, CT: Yale University Press.
Simon, H. A. (1969). *The sciences of the artificial*. Cambridge, MA: MIT Press. ((Second edition 1981))
Sloman, A. (1962). *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil, Oxford).* Retrieved from http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962
Sloman, A. (1978 Revised 2018). *The computer revolution in philosophy*. Hassocks, Sussex: Harvester Press. Retrieved from goo.gl/AJLDih
Sloman, A. (1984). The structure of the space of possible minds. In S. Torrance (Ed.), *The mind and the machine: philosophical aspects of AI.* Ellis Horwood. Retrieved from http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#49a

Sloman, A. (2002). The irrelevance of Turing machines to AI. In M. Scheutz (Ed.), *Computationalism: New Directions* (pp. 87–127). Cambridge, MA: MIT Press. Retrieved from http://www.cs.bham.ac.uk/research/cogaff/00-02.html#77

Sloman, A. (2007). *Predicting Affordance Changes.* Birmingham, UK. Retrieved from http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html

Sloman, A. (2008). Architectural and representational requirements for seeing processes, proto-affordances and affordances. In A. G. Cohn, D. C. Hogg, R. Möller, & B. Neumann (Eds.), *Logic and probability for scene interpretation.* Dagstuhl, Germany: Schloss Dagstuhl. Retrieved from http://drops.dagstuhl.de/opus/volltexte/2008/1656

Sloman, A. (2015). *What are the functions of vision? How did human language evolve?* Retrieved from http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111

Sloman, A. (2017a). Construction Kits for Biological Evolution. In S. B. Cooper & M. I. Soskova (Eds.), *The Incomputable: Journeys Beyond the Turing Barrier* (pp. 237–292). Springer International Publishing. Retrieved from http://www.springer.com/gb/book/9783319436678

Sloman, A. (2017b, Aug). *Why can't (current) machines reason like Euclid or even human toddlers?* Retrieved from http://cadia.ru.is/workshops/aga2017

Trettenbrein, P. C. (2016, Oct). The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift? *Frontiers in Systems Neuroscience*(88). Retrieved from http://doi.org/10.3389/fnsys.2016.00088

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