Chapter 1
Huge but unnoticed gaps between current AI and natural intelligence

Aaron Sloman

Abstract Despite AI's enormous practical successes, some researchers (like Turing) mainly investigate the potential of AI as science and philosophy: providing explanations of natural intelligence and answers to ancient questions about minds. Deep and difficult questions include: How could biological evolution produce multiple forms of intelligence (a diverse subset of the space of possible minds, including different stages of evolution, and different stages of development in individual organisms, functioning in widely varying contexts, etc.)? We cannot yet replicate most forms of natural intelligence. The education of AI researchers (and many others) blinds them to some of the important natural phenomena that current AI cannot model (including aspects of mathematics discussed by Kant). Current machines are not able to replicate the amazing geometrical discoveries by ancient mathematicians that remain in use worldwide, by scientists, mathematicians, engineers, architects, etc., e.g. Pythagoras' theorem. Modern geometric theorem provers start from externally provided logical axioms, based on Euclid's *Elements*, whereas for ancient mathematicians the axioms were major *discoveries*, not arbitrary starting points. Human toddlers and other animals spontaneously make similar discoveries and use them in forming intentions and planning actions – but minds vary too much to be studied through experimentally controllable phenomena. Discoveries in geometry and topology have deep, mostly unnoticed, connections with practical intelligence of many non-human species and pre-verbal human toddlers. Those discovery are unlike statistical/probabilistic learning, because, as spelled out in Kant's philosophy, they provide non-empirical knowledge of impossibilities and necessities. Analysis of current gaps between natural and artificial reasoning in topology and geometry suggests a need for future AI systems to use previously unknown forms of information processing machinery – perhaps "Super-Turing Multi-Membrane" machinery can help?

Computer Science, University of Birmingham, UK http://www.cs.bham.ac.uk/~axs,

Keywords/phrases:
AI science philosophy;
Modelling geometers?
Modelling toddlers?
Current AI gaps limitations;
Super-Turing membrane machines;
Mathematical consciousness;
Long-term AI goals.

1.1 The Meta-Morphogenesis Project

This paper opens a small window into a large project, begun over half a century ago, in my DPhil thesis [1] defending Kant's claims[2] about the nature of mathematical knowledge. After Max Clowes introduced me to AI in 1969 the project grew into an attempt to use AI to explain many aspects of minds [3], including abilities to make mathematical discoveries, especially the geometrical discoveries made by ancient mathematicians. A major new strand began in 2011, inspired by Turing's work on morphogenesis [4]. This new strand, the Meta-Morphogenesis (M-M) project,[1] investigates evolution of biological information processing mechanisms and capabilities, including an outline theory of evolved construction-kits.[2,3] This enlarged project includes the goal of understanding ancient processes and mechanisms of mathematical discovery, especially topological and geometric discovery, illustrated by the work of Archimedes, Euclid, Zeno and many others. They could not have used axiomatic, logical, forms of representation and reasoning that were not invented until recently, and which are used in current geometry theorem provers, e.g. [5].

Analysis of examples of simpler, but similar, mathematical and proto-mathematical discoveries in humans and other animals[4], suggests that intelligent animals use types of information pro-

[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] Toddler topology is illustrated in this 4.5min video (with commentary): http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/small-pencil-vid.webm

cessing 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 by Kant in 1781[2]. It is not yet clear whether virtual machines running on digital computers closely coupled with the environment (together producing something richer than a Turing machine, e.g. if the environment includes truly random quantum phenomena), are general enough. If the environment with which a digital computer interacts is not a digital machine, the coupled system, including any virtual machinery used, cannot be modelled with perfect accuracy on a Turing machine.[5]

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[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 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 understanding of natural intelligence, including For a more complete list see Margaret Boden's two-volume masterpiece [6].

---

[5] 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

1.2 Limited progress, despite spectacular successes

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. A European Commission initiative in 2004,[7] temporarily shifted the focus of robotics research in the EU back to science, but (as had happened previously in the UK Alvey Project), the initiative expanded too rapidly at a time when too few people had had the right sort of education. It also demanded practical demonstrations far too soon. As a result, the focus changed in later EU projects to demonstrable practical successes, leaving the most important scientific questions unanswered, and to some extent un-noticed!

I am not claiming that progress is impossible, only that it is very difficult and requires integration across disciplines, in a long term strategy. It also depends on a very broad and deep educational system for potential high calibre researchers. Current decisions about educational funding seem to ignore this need.

Despite the enormous practical importance of developments in AI, some researchers have always been more interested, in the potential of AI as *science* and *philosophy* than in the potential uses of AI systems. In particular AI (along with computer science) 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.

Although many science students are (unfortunately) taught to regard science as primarily concerned with finding, explaining and using observed correlations, that is a shallow view of science. Deep scientific theories all contribute to the study of what is possible and how it is possible, including the ancient atomic theory, New-

[7] Described in Colette Maloney's 2003 presentation of the scope and objectives of the EC "Cognitive Systems" initiative http://cordis.europa.eu/fp7/ict/robotics/docs/maloney-jun2003_en.pdf

ton's mechanics, chemistry, Darwin's theory of natural selection, quantum mechanics (e.g. Schrödinger's [7]), computer science and AI (as explained in Chapter 2 of [3]). The Turing-inspired Meta-Morphogenesis project mentioned in Note 1 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" [8], which I suspect is far richer than anyone currently suspects.

1.3 AI as Science and Philosophy

For most people AI is primarily an engineering activity, whereas my own interest in AI, for nearly half a century, partly inspired by Max Clowes, and works of AI founders such as Minsky e.g. [9, 10, 11], McCarthy e.g. [12, 13], and Simon, e.g. [14, 15], has centred mainly on the potential of AI to trigger and 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, 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: internal languages[8], language development, visual perception, motivational processes, and mathematical discovery). Explaining all this requires significant advances in understanding of varieties of information processing,

Clues may come from many evolutionary stages, including: microbe minds, insect minds, and evolutionary precursors of the more complex minds we hope to understand and model. This study is the Meta-Morphogenesis project mentioned in Note 1.

Unfortunately much "standard" scientific research that seeks experimental or naturally occurring regularities fails to identify

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

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. Compare how different the outputs of great composers, or poets, or novelists are, even if they live in the same location.

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

There are exceptions to these constraints, e.g. Piaget's pioneering work on childrens' understanding of Possibility and Necessity, published posthumously[16, 17]. 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 & Sauvy[18]. It could also provide useful tests for future, more human-like, robots.

1.4 Aim for generative power *vs* 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 [19]. (For historical detail see Boden [6]; Compare the claim that deep science is more concerned with discovery and explanation of *possibilities* than *laws*, in [3, 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 as I pointed out in [20]. Some preliminary suggestions regarding a "Super Turing membrane machine" are under development in Sloman [21],[9] related to ideas about affordances in [22] and McClelland's work on affordances for mental action, e.g. [23]. 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 [24].

1.4.1 An example: explaining human/animal mathematical competences

A particular generative aspect of human intelligence that has been of interest to philosophers for centuries, and discussed by Kant [2, 25], is the ability to make mathematical discoveries, including the amazing discoveries in geometry presented in Euclid's *Ele-*

[9] Compare http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html

ments 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 [21]. 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 9 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 [26] (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.)

1.4.2 AI geometry 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 version [27], and derive theorems from them 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 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 *ancient* reasoning powers. 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 by Dana Scott [28]. A simpler example, from [21], referenced in Note 9 is in Fig. 1.1.

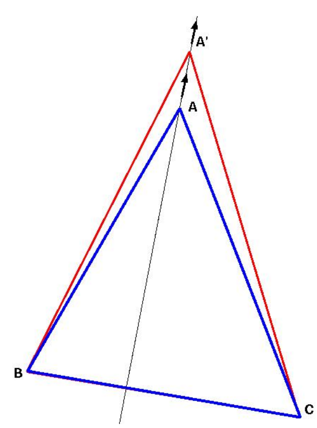

Fig. 1.1 What happens to the size of the angle at A if A is moved further from BC along a line through the opposite side BC? Notice that the answer involves thinking about two continua (the continuum of positions of the top vertex, and the continuum of angle sizes) and their relations. Many people with no mathematical training can do this easily, in my experience. What are their brains doing? How do brains represent impossibility or necessity?

If you start with an arbitrary planar triangle, like the blue one in Fig. 1.1, 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 9).

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 – a fact that is ignored by much psychological research on mathematical competences and neural theories of mathematical reasoning. This does not imply *infallibility*, as shown by Lakatos [29]. 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 [21].

1.5 Can we give robots ancient mathematical competences?

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 [4]). There is evidence in [30] that Kenneth Craik, another who died tragically young, was also thinking about such matters, perhaps inspiring Turing posthumously? Does anything in current neuroscience exppline how biological brain mechanisms can represent and reason about *perfectly straight, perfectly thin* lines, and their intersections? And reason about impossibilities, and necessary consequencs of certain kinds of motion?

Future work needs 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 [28], using diagrammatic reasoning rather than logical and arithmetic 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. How do brains enable this?

1.6 The representation of impossibility and necessity

What brain mechanisms can represent impossibility? 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?[10] 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 using semi-metrical reasoning,

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

including topological structures and partial orderings, is suggested in [31]. How can brains implement such mechanisms?

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 in Fig. 1.1. (For the effects of other orientations of the line of motion, see Note 9 and Apollonius' construction.)

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. See Note 9 and [21].

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. (Claims by Searle and others that computers cannot understand anything have been adequately refuted elsewhere. My arguments have nothing to do with his reasoning.)

1.7 Gaps in theories of consciousness

1.7.1 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 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 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 [29].

More examples of types of mathematical and non-mathematical reasoning that need to be explained and modelled are presented in references to the "*cogaff*" web site. Some discoveries of that kind seem to be made (and used) by pre-verbal human toddlers.[11]

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, as discussed by Grant [32], Trettenbrein [33], Gallistel [34] and others.

As long ago as 1944 Schrödinger[7] 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 many orders of magnitude (e.g. many centuries rather than decades). See also [35].

[11] Illustrated in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html

1.7.2 Statistical/probabilistic reasoning cannot yield 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 [36].

1.7.3 Can AI lead to robots with 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. My tentative answer is 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[37].[12] Many new biological construction kits

[12] That paper is now being extended in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

introduced new kinds of information processing mechanism, and this suggests new ideas about epigenetic processes that could produce young potential mathematicians. (Some of the ideas, about "meta-configured competences", were developed a decade ago, in collaboration with biologist Jackie Chappell [38],[13] These ideas are related to extended versions of Karmiloff-Smith's theories of "Representational Redescription" [39], and hypotheses about non-linear, structured, extendable, *internal* languages required for percepts, intentions, plans, usable generalisations, and reasoning, long before languages were used for communication[40].

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.

1.8 Other implications

This paper has opened a small window into a large complex and still growing project. There are many implications for AI as Science, AI as engineering and AI as philosophy, for which space is not available here. Neither is there space to present work in progress on requirements for a Super-Turing membrane computer, hinted at above, which may be able to replicate faithfully ancient mathematical discovery processes and related discovery processes in pre-verbal humans and many other intelligent species. The Meta-Morphogenesis web site already includes some promising ideas and is expected to continue growing.[14]

---

[13] Currently being elaborated in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html

[14] See Note 1, and also
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-geom.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

References

1. Sloman, A.: Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil Thesis). PhD thesis, Oxford University (1962) http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962.
2. Kant, I.: Critique of Pure Reason. Macmillan, London (1781) Translated (1929) by Norman Kemp Smith.
3. Sloman, A.: The Computer Revolution in Philosophy. Harvester Press (and Humanities Press), Hassocks, Sussex (1978) http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp, Revised 2015.
4. Turing, A.M.: The Chemical Basis Of Morphogenesis. Phil. Trans. R. Soc. London B 237 **237** (1952) 37–72 http://dx.doi.org/10.1098/rstb.1952.0012.
5. Chou, S.C., Gao, X.S., Zhang, J.Z.: Machine Proofs In Geometry: Automated Production of Readable Proofs for Geometry Theorems. World Scientific, Singapore (1994) http://www.mmrc.iss.ac.cn/ xgao/paper/book-area.pdf.
6. Boden, M.A.: Mind As Machine: A history of Cognitive Science (Vols 1–2). Oxford University Press, Oxford (2006)
7. Schrödinger, E.: What is life? CUP, Cambridge (1944)
8. Sloman, A.: The structure of the space of possible minds. In Torrance, S., ed.: The Mind and the Machine: philosophical aspects of Artificial Intelligence. Ellis Horwood, Chichester (1984) http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#49a.
9. Minsky, M.L.: Steps toward artificial intelligence. In Feigenbaum, E., Feldman, J., eds.: Computers and Thought. McGraw-Hill, New York (1963) 406–450
10. Minsky, M.L.: Matter Mind and Models. In Minsky, M.L., ed.: Semantic Information Processing. MIT Press, Cambridge, MA (1968)
11. Minsky, M.L.: The Emotion Machine. Pantheon, New York (2006)
12. McCarthy, J.: Ascribing mental qualities to machines. In Ringle, M., ed.: Philosophical Perspectives in Artificial Intelligence. Humanities Press, Atlantic Highlands, NJ (1979) 161–195 http://www-formal.stanford.edu/jmc/ascribing/ascribing.html.
13. McCarthy, J.: The well-designed child. Artificial Intelligence **172**(18) (2008) 2003–2014 http://www-formal.stanford.edu/jmc/child.html.
14. Simon, H.A.: Motivational and emotional controls of cognition. In Simon, H.A., ed.: reprinted in Models of Thought. Yale University Press, Newhaven, CT (1967) 29–38
15. Simon, H.A.: The Sciences of the Artificial. MIT Press, Cambridge, MA (1969) (Second edition 1981).
16. Piaget, J.: Possibility and Necessity Vol 1. The role of possibility in cognitive development. U. of Minnesota Press, Minneapolis (1981) Tr. by Helga Feider from French in 1987.
17. Piaget, J.: Possibility and Necessity Vol 2. The role of necessity in cognitive development. U. of Minnesota Press, Minneapolis (1983) Tr. by Helga Feider from French in 1987.
18. Sauvy, J., Sauvy, S.: 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.
19. Chomsky, N.: Aspects of the theory of syntax. MIT Press, Cambridge, MA (1965)
20. Sloman, A.: The irrelevance of Turing machines to AI. In Scheutz, M., ed.: Computationalism: New Directions. MIT Press, Cambridge, MA (2002) 87–127 http://www.cs.bham.ac.uk/research/cogaff/00-02.html#77.
21. Sloman, A.: Why can't (current) machines reason like Euclid or even human toddlers? (And many other intelligent animals) (Aug 2017) http://cadia.ru.is/workshops/aga2017.

22. Sloman, A.: Architectural and representational requirements for seeing processes, proto-affordances and affordances. In Cohn, A.G., Hogg, D.C., Möller, R., Neumann, B., eds.: Logic and Probability for Scene Interpretation. Number 08091 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany (2008)
23. McClelland, T.: AI and affordances for mental action. In: Computing and Philosophy Symposium, Proceedings of the AISB Annual Convention 2017, Bath (Apr 2017) 372–37 http://wrap.warwick.ac.uk/87246.
24. Senghas, A.: Language Emergence: Clues from a New Bedouin Sign Language. Current Biology **15**(12) (2005) R463–R465 http://dx.doi.org/10.1016/j.cub.2005.06.018.
25. Kant, I.: Prolegomena to Any Future Metaphysics That Will Be Able to Present Itself as a Science. (1783) Originally Prolegomena zu einer jeden kunstigen Metaphysik, die als Wissenschaft wird auftreten konnen. Several translated versions online.
26. Penrose, R.: Shadows of the mind: A Search for the Missing Science of Consciousness. OUP, Oxford (1994)
27. Hilbert, D.: The Foundations of Geometry. Project Gutenberg, Salt Lake City (1899) Translated 1902 by E.J. Townsend, from 1899 German edition, http://www.gutenberg.org/ebooks/17384.
28. Scott, D.S.: Geometry without points (2014) https://www.logic.at/latd2014/2014
29. Lakatos, I.: Proofs and Refutations. Cambridge University Press, Cambridge, UK (1976)
30. Craik, K.: The Nature of Explanation. Cambridge University Press, London, New York (1943)
31. Sloman, A.: Predicting Affordance Changes (Alternative ways to deal with uncertainty). Technical Report COSY-DP-0702, School of Computer Science, University of Birmingham, Birmingham, UK (Nov 2007) http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html.
32. Grant, S.G.: Computing behaviour in complex synapses: Synapse proteome complexity and the evolution of behaviour and disease. Biochemist **32** (April 2010) 6–9 http://www.biochemist.org/bio/default.htm?VOL=32&ISSUE=2.
33. Trettenbrein, P.C.: The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift? Frontiers in Systems Neuroscience (88) (Oct 2016) http://doi.org/10.3389/fnsys.2016.00088.
34. Gallistel, C., Matzel, L.: The neuroscience of learning: beyond the Hebbian synapse. Annual Revue of Psychology **64** (2012(Epub)) 169–200 DOI: 10.1146/annurev-psych-113011-143807.
35. Newport, T.: Brains and Computers: Amino Acids versus Transistors. (2015) https://www.amazon.com/dp/B00OQFN6LA.
36. Chappell, J., Sloman, A.: Contributions to WONAC: International Workshop on Natural and Artificial Cognition Two ways of understanding causation: Humean and Kantian (June 25-26 2007) http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac.
37. Sloman, A.: Construction Kits for Biological Evolution. In Cooper, S.B., Soskova, M.I., eds.: The Incomputable: Journeys Beyond the Turing Barrier. Springer International Publishing (2017) 237–292 http://www.springer.com/gb/book/9783319436678.
38. Chappell, J., Sloman, A.: Natural and artificial meta-configured altricial information-processing systems. International Journal of Unconventional Computing **3**(3) (2007) 211–239 http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717.
39. Karmiloff-Smith, A.: Beyond Modularity: A Developmental Perspective on Cognitive Science. MIT Press, Cambridge, MA (1992)
40. Sloman, A.: What are the functions of vision? How did human language evolve? (2015) http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111.