School of Computer Science | University of Birmingham UK

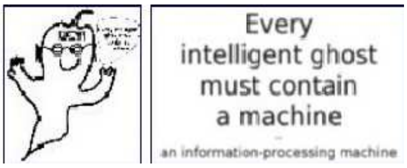

Departmental Seminar, School of Computer Science
The University of Birmingham

Why is it so hard to get machines to reason like our ancestors who produced Euclidean Geometry?

Speaker: Aaron Sloman

4pm Thursday 20th June 2013,
Official abstract for the talk:
http://talks.bham.ac.uk/talk/index/1063
This informal, extended abstract may be modified from time to time.
Last modified: 18 Jun 2013; 7 Dec 2014

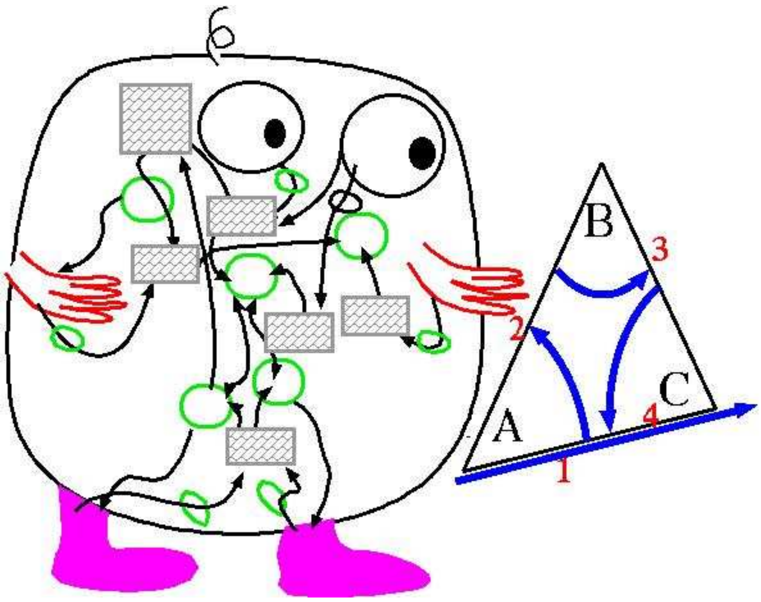

Extended Abstract

The ideas presented here cannot be presented in a clear and simple cumulative fashion: there’s too much mutual dependence between sub-topics that do not naturally form a linear order. So I have to hope for readers who are persistent and patient, as well as highly critical.

The universe seems to include (potentially) infinitely many domains, of many different sorts, with many different degrees and kinds of complexity. The domains instantiated on this planet seem to have increased steadily in size and complexity and diversity since life began, partly because of the

increasing diversity and complexity of forms of life, and their interactions.

A domain is a set of possible **structures** and possible **processes** involving those structures. For quick introductions to many domains see Vi Hart's amazing high speed online math-doodles.

The fact that a specific type of structure or process is possible in a domain does not imply that instances of that specific type will ever exist. For example, it is likely that there are hugely many 500-word possible (grammatical, meaningful, sentences of English that will never be uttered, and many possible machines and buildings that will never be constructed. The domain of shape changes of a 5 metre long piece of string lying on a flat surface may include process-types that will never be instantiated. It's even more likely that the domain of possible shape transitions of a climbing plant, such as ivy, is so vast that despite the many millions of actual instances of ivy growing, only a small subset of the possibilities will ever be realised. (To some extent this will depend on how demanding the criteria are for saying that two slightly different instances are of the same type.)

Both the instances included in many of the domains, and the domains themselves have become bigger, more complex, and more varied over time -- for example, as more species of organisms have evolved and, the size and complexity of organisms and their capabilities have increased.

As organisms acquired more complex sensors, effectors, and control mechanisms, including mechanisms that change their behaviour through learning, and mechanisms that take account of information processing and control processes in other organisms, increasing varieties of domains of information processing have developed including domains of information about information, both in individual meta-cognition, and in cognition concerned with information-processing (sensing, learning, wanting, planning, choosing, acting on intentions, hypothesising, reasoning, etc.) in others.

A major driver of this process has been biological evolution, perhaps the richest positive feedback system known.

Each domain comprises a set of possible **structures** and **processes** (transformations), in which structures change.

Some domains are very simple, (e.g. the domain of possible ways a coin tossed once can land: heads up or tails up). Others are more complex, including

- discrete domains

  (e.g. sequences of coin-tosses, truth-tables, the domain of initial segments of the natural number series: 1, 1-2, 1-2-3, 1-2-3-4, etc, or the domain of counting processes, setting up correspondences between such segments and collections of other entities),

- continuous domains (e.g. motions of lines in a plane, motions of a flexible string),
- mixed continuous/discrete domains (e.g. various biological domains including chemical structures and processes, or the domain of Lego construction processes),
- meta-domains (e.g. the domain of types of algebraic domains, the domain of types of domains of grammatical sentences, the domain of types of domains of programs expressible in a particular programming language).

A more detailed account of domains and their relationships to mathematics is under development here.

Organisms need repertoires of behaviour related to their structures, their capabilities, their varying needs and varying circumstances. Domains of possible behaviours for organisms become increasingly complex and varied as evolution produces more complex and varied organisms that interact with other organisms (as prey, as predators, as mates, as offspring, as parents, as collaborators, as competitors, as symbionts) as well as with changing physical circumstances produced by weather, earthquakes, volcanoes, floods, migration, etc.

Some domains of possible behaviours are produced (very slowly) by biological evolution. Later, evolution produced organisms that can discover domains relevant to their capabilities and circumstances and produce appropriate reactions to particular circumstances, for instance, using learnt reflexes and responses to triggers of various sorts. (Varieties of online intelligence.)

Later, as the challenges and opportunities became more complex, organisms evolved, and learnt/developed more sophisticated ways of selecting behaviours by reasoning about sets of possibilities. For this they required information about the domains of possibilities and constraints on possibilities. This required evolution to produce new abilities to acquire knowledge about how to use information about domains, for example in constructing plans, finding explanations for perceived situations or events, etc. (Varieties of offline intelligence.)

The elements of some domains can be physically instantiated, whereas others cannot, though for different reasons, e.g. the domain of impossible 3-D objects depictable in 2-D line drawings, as in Escher's pictures -- impossible to instantiate because spatial constraints would be violated -- or the domain of possible re-orderings of the set of natural numbers -- impossible to instantiate if the universe is finite.) A major challenge for biology, psychology, neuroscience, AI/Robotics and philosophy is to identify the various kinds of domain, with various origins, that have proved relevant in one way or another to individuals, to species, to social groups, to ecosystems, to nations, ..., and to explain how individuals can come to know about and use information about such domains, including, in some cases creating new domains with greater power (e.g. adding metrics to pre-existing spatial domains, inventing calculus, inventing programmable looms, etc.).

It seems that without a certain type of richness in the domain of possible chemical structures and processes there would not have been life as we know it (Ref T. Ganti).

Each form of life has a domain of possible sensory-motor interactions with each of the types of environment in which that life is possible. More complex forms of interaction became possible when evolution produced organisms that could explore and learn about domains and deploy that knowledge, e.g. in planning action sequences, or building maps of extended spatial structures on the basis of sequences of sensory motor interactions. Using such domain knowledge enabled more intelligent animals to detect and reason about positive and negative affordances for action. (Ref: J.J. Gibson).

In humans some of that led to meta-knowledge about such competences and eventually to the production of Euclid's **Elements**, one of the greatest achievements of biological evolution. Many other forms of mathematical knowledge grew out of later explorations of domains, and then meta-domains of many kinds.

One of the key features of such knowledge is that it concerns grasping some set of **possibilities** and then discovering **constraints** on those possibilities, e.g. learning that some extensions of a set of possibilities can be described or depicted but are not included in the set, e.g. the penrose triangle or a set of 3 objects combined with a non-overlapping set of 2 objects forming a set of 4 objects.

This is totally different from and more fundamental than discovery of probabilities through empirical observations, the current focus of huge amounts of research in AI/Robotics and neuroscience (much of it misguided in my opinion).

In the last few decades there have been tremendous advances in AI theorem proving techniques, and we now have programs that can find and prove theorems that would defeat most humans, including a package that will sell you a new, unique, non-trivial theorem named after you (REF). But it has proved extremely difficult to get computers to engage in the kinds of reasoning even a human toddler can do and some other animals seem able to do that made the development of Euclidean geometry possible. This talk will present some examples and discuss possible ways of making progress, with potential implications for developmental psychology, neuroscience, theories of animal cognition, and philosophy of mathematics, as well as AI and Robotics.

This abstract is available at:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/euclidean-ancestors.html

A partial index of discussion notes in this directory is in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html

Examples of simple forms of geometrical reasoning that seem hard to model on computers are in:
http://tinyurl.com/CogMisc/triangle-theorem.html
Hidden Depths of Triangle Qualia

This is part of the Meta-Morphogenesis project:
http://tinyurl.com/CogMisc/meta-morphogenesis.html

The ideas here are closely related to some of the ideas developed by Annette Karmiloff-Smith in her 1992 book: Beyond Modularity, briefly reviewed here:
http://tinyurl.com/CogMisc/beyond-modularity.html

A partial survey of types of "toddler theorems" is here (contributions welcome):
http://tinyurl.com/CogMisc/toddler-theorems.html

To be modified and extended.

A video interview introducing some of these topics:
http://www.youtube.com/all_comments?v=iuH8dC7Snno

Notes: Euclidean geometry is one of the greatest products of human minds, brought together in Euclid's **Elements** over two millennia ago.

However, at some distant earlier time there were no geometry textbooks and no teachers. So, long before Euclid, our ancestors, perhaps while building huts, temples and pyramids, measuring fields, making tools or weapons, or reasoning about routes, must have noticed facts about spatial structures and processes that are both useful (like facts about physics, geography, biology, and human languages), but are also demonstrable by reasoning with logic and diagrams. Mathematicians do not have to keep checking that their discoveries remain true at high altitudes, or in cold weather, or on surfaces with unusual materials or colours -- because they can **prove** things.

Without teachers to help, biological evolution must somehow have produced information-processing mechanisms that allowed ancient humans to develop the concepts, notice the relationships and discover the proofs that their descendants are taught at school, but which we have the ability to discover for ourselves, as our ancestors did.

This suggests that normal human children have the potential to make those discoveries themselves, under appropriate conditions. I suspect there are also deep connections with competences that have evolved in other intelligent species that understand spatial structures, relationships and processes -- such as nest-building birds, squirrels that steal nuts from bird feeders, elephants that manipulate water, mud, sand and foliage with their trunks, and apes coping with many complex structures as they move through and feed in tree-tops.

Can we replicate evolution's achievements, and create robots that start off with competences of young children and later, as they develop, make simple discoveries in Euclidean geometry? I'll explain why that's hard to do -- but perhaps not impossible. There have been great advances getting computers to reason logically, algebraically and arithmetically, but the kinds of reasoning in Euclid, e.g. using diagrams, are very different.

I'll discuss some of the problems and possible ways forward. Perhaps, someone now studying geometry and computing at school will one day design the first baby robot that grows up to be a self-taught robot geometer, and, like some of our ancestors, discovers for itself why the angles of a triangle must add up to exactly half a rotation.

Other related materials

- http://vihart.com/
  Vi Hart's brilliant online high-speed mathematical doodling videos introduce many different domains. This web page gives a small selection: http://vihart.com/doodling/

- http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html
  Biology, Mathematics, Philosophy, and Evolution of Information Processing
  Mathematics is at root a biological, not anthropological, phenomenon.
  (Under development.)

- http://math.berkeley.edu/~rbayer/09su-55/handouts/ProofByPicture-printable.pdf
  Robertson Bayer, Proof By Picture (PDF lecture slides),
  University of California, Berkeley Math 55, Summer 2009
  (A collection of diagrammatic proofs of mathematical theorems, most of them non-geometric -- e.g. geometric proofs of theorems in number theory.
  Includes the 'Chinese' proof of Pythagoras' Theorem.)

- http://uhra.herts.ac.uk/dspace/bitstream/2299/2455/1/phd11thjulyformattedprinter.pdf
  Cathal Butler
  Evaluating the Utility and Validity of the Representational Redescription Model as a General Model for Cognitive Development
  PhD thesis, University of Hertfordshire, October 2007

- Kenneth Craik, **The Nature of Explanation,**
  Cambridge University Press, 1943, London, New York,

  NOTE:
  Craik proposed that biological evolution produced animals with the ability to work out what the consequences of an action would be without performing the action, by making use of an abstract **model** of the situation in which the action is performed. It is not clear to me that he noticed the difference between running a detailed model of a specific situation to discover the specific consequences, which some current AI systems (e.g. game-engines) can do, and noticing an invariant property of such a process with different starting configurations as required for understanding why a strategy will work in a (possibly infinite) class of cases.
  I think he came close, but did not quite get there, but I have read only the 1943 book.
  Compare:
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html

- George B. Dyson,
  Darwin Among The Machines: The Evolution Of Global Intelligence,
  Addison-Wesley, 1997,

- **Entertaining video tutorial on Pythagoras and irrationality**
  showing the connection between irrationality of square root of 2, and Pythagoras' theorem.
  https://www.youtube.com/watch?v=X1E7I7_r3Cw
  "What's up with Pythagoras", by Vi Hart.
  See more of her work: https://www.youtube.com/user/Vihart
  And here: http://vihart.com
  You may have to pause and replay bits of her videos, to take in some of the details.

- Tibor Ganti, *The Principles of Life,*
  Eds. E\ors Szathm\'ary and James Griesemer,
  Translation of the 1971 Hungarian edition, with notes, OUP, New York, 2003,
  Usefully summarised in http://wasdarwinwrong.com/korthof66.htm

- http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#gibson
  Talk 93: What's vision for, and how does it work?
  From Marr (and earlier) to Gibson and Beyond

- Immanuel Kant, **Critique of Pure Reason,** 1781,
  Translated (1929) by Norman Kemp Smith, London, Macmillan,

- Annette Karmiloff-Smith,
  Beyond Modularity: A Developmental Perspective on Cognitive Science,
  MIT Press, 1992,
  Summarised in: http://tinyurl.com/CogMisc/beyond-modularity.html

- http://adrenaline.ucsd.edu/Kirsh/Articles/Earwig/earwig-cleaned.html
  David Kirsh, "Today the earwig, tomorrow man?,
  in **Artificial Intelligence,** 47, 1, pp. 161--184, 1991

- Imre Lakatos, **Proofs and Refutations,**
  1976, CUP, Cambridge, UK,

- Roger B. Nelsen, (Often mis-spellt as "Nelson")
  Proofs without words: Exercises in Visual Thinking,
  Mathematical Association of America, Washington DC, 1993,

- http://hdl.handle.net/10289/5153
  Catherine Legg, What is a logical diagram?
  Paper presented at Mini-Conference on Logical Pragmatism,
  Auckland, New Zealand, February 25, 2011.

- **Mary Leng interviews Mateja Jamnik on Spatial Reasoning**,
  In **The Reasoner** 7(1), Jan 2012. pp 1--4.
  http://www.kent.ac.uk/secl/philosophy/jw/TheReasoner/vol7/TheReasoner-7(1).pdf

- John McCarthy, Ascribing mental qualities to machines,
  in *Philosophical Perspectives in Artificial Intelligence,*
  Ed. M. Ringle, Humanities Press, 1979, pp.161--195,
  http://www-formal.stanford.edu/jmc/ascribing/ascribing.html

- J. Sauvy and S. Sauvy,
  The Child's Discovery of Space: From hopscotch to mazes --
  an introduction to intuitive topology,
  Penguin Education, 1974, Translated from the French by Pam Wells,

- http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#706
  Aaron Sloman,
  Knowing and Understanding: Relations between meaning and truth,
  meaning and necessary truth, meaning and synthetic necessary truth
  DPhil Thesis, Oxford University, 1962,

- http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#701
  Aaron Sloman, 'Necessary', 'A Priori' and 'Analytic',
  Analysis, 26, 1, 1965, pp. 12--16,

- http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#714
  Aaron Sloman, Functions and Rogators,
  In *Formal Systems and Recursive Functions:*
  Proceedings of the Eighth Logic Colloquium Oxford, July 1963,
  Eds. J. N. Crossley and M. A. E. Dummett,
  North-Holland, Amsterdam, 1965, pp. 156--175,

- http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#712
  Aaron Sloman, Explaining Logical Necessity,
  In *Proceedings of the Aristotelian Society,*
  vol 69 1968/9, pp. 33--50,

- http://tinyurl.com/BhamCog/04.html#200407
  Aaron Sloman, 'Interactions between philosophy and AI:
  The role of intuition and non-logical reasoning in intelligence',
  Proc 2nd IJCAI, 1971, London, pp. 209--226, William Kaufmann,
  Also reprinted in AI journal, 1971,
  and in "The Computer Revolution in Philosophy", Chapter 7, 1978)
  http://www.cs.bham.ac.uk/research/projects/cogaff/crp#chap7
  That article was originally written as a critique of McCarthy and Hayes 1969:
  "Some philosophical problems from the standpoint of Artificial Intelligence"
  (In **Machine Intelligence 4**)

- http://tinyurl.com/BhamCog/96-99.html#15
  Aaron Sloman, Actual Possibilities, in
  Principles of Knowledge Representation and Reasoning: Proc. 5th Int. Conf. (KR '96),
  Eds. L.C. Aiello and S.C. Shapiro, Morgan Kaufmann, Boston, MA, 1996, pp. 627--638,

- http://dspace.mit.edu/handle/1721.1/6894
  Gerald. J. Sussman,
  A computational model of skill acquisition, American Elsevier, 1975,

- http://theorymine.co.uk
  Personalized mathematical theorems

- http://www.cns.atr.jp/~emre/papers/PhDThesis.pdf
  Emre Ugur, A Developmental Framework for Learning Affordances (PhD thesis),
  Graduate School of Natural and Applied Sciences,
  Middle East Technical University, Ankara, Turkey, 2010,

- Max Wertheimer, **Productive Thinking**
  New York, Harper and Brothers, 1945

Installed: 11 Jun 2013
Last updated: 12 Jun 2013; 30 Oct 2020: Formatting only

Maintained by Aaron Sloman
School of Computer Science
The University of Birmingham