Every
intelligent ghost
must contain
a machine
–
an information-processing machine

Gaps between biological and artificial intelligence:
A high-level overview

Expanded background material for the tutorial on the Meta-Morphogenesis project, presented on Sunday 10th July 2016.

A more recent version, including a video presentation was prepared for an invited talk at a workshop at IJCAI 2017
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html
Why can't (current) machines reason like Euclid
or even human toddlers?
(And many other intelligent animals)
Video presentation here(42 min)
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17

(DRAFT: Liable to change)

Aaron Sloman
School of Computer Science, University of Birmingham
(Philosopher in a Computer Science department)

Installed: 6 Aug 2016
Last updated: 9 Sep 2017 (Link to 2017 talk+video)
This paper is
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-ai-gaps.html
A PDF version may be added later.

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

Background to this document.

This web location originally held a submission to the IJCAI 2016 workshop on
Bridging the Gap between Human and Automated Reasoning
http://ratiolog.uni-koblenz.de/bridging2016 held at the International Joint Conference on AI,
New York, July 2016: http://ijcai-16.org/

The submission entitled "Natural Vision and Mathematics: Seeing Impossibilities" was accepted,
and a revised version is in the workshop proceedings, and also available at:
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-bridging-gap-2016.pdf

This is now a new but related paper, providing background for my tutorial, also presented at IJCAI
2016 (10th July 2016, New York):

> Tutorial T24: If Turing had lived longer, how might he
> have investigated what AI and Philosophy can learn
> from evolved information processing systems?
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-tut-ijcai-2016.html

The tutorial presented a small sample of features of natural intelligence (in non-human animals,
pre-verbal humans and adult humans, e.g. Euclid, Archimedes) that are not matched in current
AI/Robotic systems. Those gaps are usually not noticed by AI researchers, especially if they are
mainly concerned with solving narrowly defined practical problems that ignore most aspects of
natural intelligence.

The tutorial referred to the Meta-Morphogenesis (M-M) project, which has the aim of identifying
many such gaps and attempting to relate them to varieties of information processing produced by
biological evolution. I conjecture that there are many ways in which AI falls short of natural
intelligence. Moreover, it is not clear whether that's simply because the requirements have not
been noticed by AI researchers or in part because the forms of information processing that are
currently deployed in AI are inadequate for the purpose of replicating natural intelligence, apart
from special subsets, or simply because researchers have not yet understood the problems well
enough to produce appropriate designs.

The M-M project aims to identify both unnoticed types of information-processing competence tha
exist in animals and also unnoticed types of information-processing mechanism used by those
competences. By surveying familiar and unfamiliar transitions in information processing
competences and mechanisms since the earliest life forms we may discover intermediate types of
competence and mechanism that have not yet been noticed, some of which still play important
roles in human brains and minds.

Since time limits (and limitations of current knowledge) restricted the range of examples presented
in the tutorial, this document provides a much larger collection of examples -- with corresponding
challenges for AI researchers, neuroscientists, psychologists and philosophers. But this is merely a
draft, incomplete, list of types of biological information processing to illustrate the scope of what is
already known (by some researchers, though not all). Over time I expect to extend this list and
extend the pointers to more detailed research documents exploring the competences and
mechanisms required.

What follows is a list of AI-Gaps (AIGs) many of which have not been noticed, or have been ignored, by researchers in AI, cognitive science, psychology and philosophy. I think filling these gaps with previously unknown information-processing mechanisms is a pre-requisite for achieving many of the long term scientific (explanatory) and engineering (useful) goals of AI.

Although several of the founders of AI, mentioned above, were interested in AI as science, some of them (implicitly or explicitly) restricted their research to the science of human-like intelligent systems. I suspect the aim of understanding late products of biological evolution while ignoring most of the intermediate, less complex cases, may be an impossible task. I think McCarthy implicitly acknowledged this in his 1996 paper "The Well Designed Child" [McCarthy Child].

Summary of gaps between current artificial and biological intelligence

There are many deep, but not widely recognized, gaps between natural intelligence (NI) (in humans and non-humans) and products of AI research and development. Examples of such gaps are given below. The gaps may not matter for narrowly focused AI applications (AI as engineering).

However they do matter for AI construed as the science of intelligence, aiming to use computational theories to answer deep questions about the variety of forms of intelligence, including intelligence in humans and other animals. Several of the founders of AI, including Turing, McCarthy, Minsky, Newell and Simon were not concerned only with practical applications: they hoped to use computational concepts, theories and techniques developed in AI to understand and model aspects of natural intelligence, though for a time some of them underestimated the difficulties.

Now in 2016, 60 years after the Dartmouth conference, huge gaps between natural and artificial intelligence remain, but to many AI researchers, and their admirers and critics the gaps are invisible, partly because of the spectacular successes of AI in tasks that previously seemed to need natural intelligence based on biological brains, and partly because there are many aspects of human and animal intelligence that go unnoticed, perhaps because they are too familiar to seem deep and difficult to explain. Piaget was unusual in paying attention to such aspects, but sought explanatory models based on hopelessly inadequate concepts and tools. Long before him Immanuel Kant had noticed some of the problems [Kant 1781], but lacked the computational concepts and theories we now have.

However, the computational tools, techniques, concepts and theories now available have not yet been shown to be sufficient to model or explain all aspects of natural intelligence (NI) -- important gaps remain, some of them presented below and in the tutorial. However, biological evolution based on physics and chemistry produced NI, so there must be things we have not yet understood about how that was achieved, and how the physical world originally made that possible. If we understood how, then perhaps we could use that understanding as a basis for AI systems that come closer the major achievements of natural intelligence.

But that requires researchers to understand what has been achieved: it is normally assumed that philosophers, psychologists, linguists, neuroscientists, and biologists studying other species can tell us about the capabilities of humans and other animals, and therefore what needs to be explained. But the history of science shows that many of phenomena that need to be explained are invisible to those who have not had the experience of developing and testing apparently successful explanatory theories (e.g. the Ptolemaic theory of planetary motion, Newton's mechanics) and then

finding their gaps and errors.

Likewise, many aspects of human and animal intelligence will remain invisible to those who have not developed and tested, or at least become familiar with, powerful and initially plausible theories, and then observed where they fail. A consequence of that invisibility is that shallow and inadequate explanatory ideas can become fashionable. (Fashionable embodiment-based theories were criticised in [Sloman,2009]. See also [Rescorla 2015].)

The Meta-morphogenesis project (partly inspired by Turing's work on morphogenesis) is based on a conjecture that one way to find clues regarding how to bridge those gaps in current AI (and neuroscience) is to try to identify previously unnoticed transitions in biological evolution: including all major transitions in forms of information processing between the very simplest organisms (or pre-life molecules) and the most sophisticated existing life forms. Each transition produces one or more of the following (not intended to be a complete list):

- New uses of information
- New sources of information
- New forms of information
- New transformations of information
- New forms of representation of information
- New information detectors (sensors)
- New mechanisms for storing, manipulating, transforming and using information
- New controlled physical behaviours
- New forms of learning and development
- New types of genetic information
- New reproductive mechanisms
- New epigenetic mechanisms
- New or extended ontologies
- New information processing architectures, with multiple interacting subsystems
- New forms of interaction between information-processing subsystems
- New distributed types of information processing (distributed across different individuals, or between agents and their environments)
- New forms of meta-cognition, meta-meta-cognition, meta-meta-...
- New control strategies (including new forms of meta-management)
- New values, preferences, goals, ideals, standards, ...
- New forms of communication
- New uses for communication
- New forms of collaboration
- New shared forms of information processing (cultural mechanisms)
  (Not necessarily all collaborative)
- New construction kits (discussed in more detail in
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

Some of the products of evolution include side-effects that alter the mechanisms of evolution: hence the label "Meta-morphogenesis" for the project.

Although the concept of "information" is both very old and very widely used (e.g. by Jane Austen in *Pride and Prejudice* a hundred years before Shannon [Austen(Information)] it is frequently misrepresented as being concerned with transmission and storage of messages, whereas information is important because it can be **used**. Evolution continually discovered new uses for information and new types of information, new information representations and new information-processing mechanisms. Filling gaps in our knowledge about this may provide important new clues for AI. This concept of information is seriously distorted by the common assumption (especially following Shannon) that information must have a numerical measure.

Research triggered by such gaps may draw attention to previously unnoticed intermediate evolved information processing abilities and mechanisms, and uses of information. Some of those previously unnoticed uses and mechanisms evolved long ago may still perform important functions in human brains -- functions that have not been noticed e.g. because current brain research methods cannot identify the functions or the mechanisms.

In particular, the emphasis on uses of information has begun to draw attention to the diversity of forms of representation, i.e. languages produced by evolution (or by individual development guided partly by the environment and partly by the genone), including languages for purely **internal** uses in perception, intending, wondering, noticing, discovering, planning, deciding, carrying out plans, and many more. If these occur in many intelligent non-human animals, that must **completely** revised our view of the nature of language (illustrated below and in [Sloman(Vision)]).

Epigenesis: Attending to previously unnoticed transitions in cognitive development of individuals in intelligent human and non-human species may also provide clues (e.g. use of topological information by pre-verbal human toddlers and other animals)[Chappell Sloman 2007], [Karmiloff-Smith 1992].

Like McCarthy and Minsky, I focus more on AI as **science** and **philosophy** than AI as **engineering**, though the interests overlap: many aspects of AI as engineering depend on good science and philosophy. In particular, researchers in AI (and cognitive science) who know nothing about the work of Kant and other great philosophers risk missing some of the deepest features of minds, language, and thought that need to be explained and modelled.

The reverse is also true: philosophers with shallow understanding of the science and engineering issues in computing and AI, including what we have learnt about varieties of virtual machinery since Turing died, will produce shallow philosophical theories of mind, language, science, mathematics, etc.

Many researchers remain mystified, or even mystical, about mental phenomena because their education has not introduced them to the required types of explanatory mechanism -- mechanisms capable of filling the so-called "Explanatory Gap" (pointed out by Darwin's admirer T.H.Huxley [Huxley 1866/1872] and repeatedly re-discovered, and re-labelled, since then)[SEP "Consciousness"]. (Huxley toned down his wording in the 1872 edition.)

Acknowledgements:
This paper owes much to discussions with Jackie Chappell about animal intelligence.

References

[Austen(Information)]
Jane Austen's concept of information (As opposed to Claude Shannon's) (Online discussion note.)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html

[Beaudoin, 2014]
L.P. Beaudoin, 2014, *Cognitive Productivity: The Art and Science of Using Knowledge to Become Profoundly Effective,* Leanpub. http://leanpub.com/cognitiveproductivity

[Beaudoin, CogZest]
Luc Beaudoin (2015??), *The CogZest web site* https://cogzest.com/
https://cogzest.com/2015/07/a-tale-of-two-summer-conferences-isre-2015-and-cogsci-2015/

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

[Crick, 1954]
F. H. C. Crick, 1954/2015 The structure of the hereditary material, in *Nobel Prizewinners who changed out world* Scientific American, Topix Media Lab, New York USA 1954/2015 pp. 6--15

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

[Clowes 1971]
Clowes, M. (1971). On seeing things. Artificial Intelligence, 2 (1), 79-116.

[Cohen 1962]
L.J. Cohen, 1962, *The diversity of meaning,* Methuen \& Co Ltd, London

[Dennett 1995]
Dennett, D. (1995). Darwin's dangerous idea: Evolution and the meanings of life. London and New York: Penguin Press.

[Frege 1950]
Frege, G. (1950). The Foundations of Arithmetic: a logico-mathematical enquiry into the concept of number. Oxford: B.H. Blackwell. ((Tr. J.L. Austin. Original 1884))

[Ganti,Life]
Tibor Ganti, 2003. *The Principles of Life,*
Eds. E. Szathmáry, & J. Griesemer, (Translation of the 1971 Hungarian edition), OUP, New York.
See the very useful summary/review of this book by Gert Korthof:
http://wasdarwinwrong.com/korthof66.htm

[Gibson 1979]
Gibson, J J. (1979). *The ecological approach to visual perception.* Boston, MA: Houghton Mifflin.

[Glasgow, Narayanan, Chandrasekaran; . 1995]
Glasgow, J., Narayanan, H. Chandrasekaran, B. (). (Eds.). (1995). *Diagrammatic reasoning: Computational and cognitive perspectives.* Cambridge, MA: MIT Press.

[Jablonka Lamb 2005]
Jablonka, E. Lamb, M J. (2005). Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life. Cambridge MA: MIT Press.

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

[Lakatos 1976]
    Lakatos, I. (1976). Proofs and Refutations. Cambridge, UK: Cambridge University Press.

[Marr 1982]
    Marr, D. (1982). Vision. San Francisco: W.H.Freeman.

[McCarthy_HayesMcCarthy_Hayes 1969]
    McCarthy, J. Hayes, P. (1969). Some philosophical problems from the standpoint of AI. In B. Meltzer & D. Michie (Eds.), Machine Intelligence 4 (pp. 463-502). Edinburgh, Scotland: Edinburgh University Press.

[McCarthy Child]
    John McCarthy (1996). "The Well Designed Child" (unpublished research paper)
    http://www-formal.stanford.edu/jmc/child.html
    Later published in the *AI Journal*, 172, 18, pp 2003--2014, 2008

[Piaget 1952]
    Jean Piaget, (1952). *The Child's Conception of Number.* London: Routledge & Kegan Paul.

[Piaget 1981-1983]
    Jean Piaget (1981,1983) *Possibility and Necessity* (1981/1983) *Vol 1. The role of possibility in cognitive development* (1981), *Vol 2. The role of necessity in cognitive development* (1983), (Tr. by Helga Feider from French in 1987)

[Rescorla 2015]
    Michael Rescorla, (2015), The Computational Theory of Mind, in *The Stanford Encyclopedia of Philosophy* Ed. Edward N. Zalta, Winter 2015,
    http://plato.stanford.edu/archives/win2015/entries/computational-mind/

[Senghas 2005]
    Senghas, A. (2005). Language Emergence: Clues from a New Bedouin Sign Language. Current Biology, 15 (12), R463-R465.

[Sloman 1962]
    Sloman, A. (1962). Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil Thesis).
    http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962

[Sloman 1971]
    Sloman, A., (1971). Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In **Proc 2nd IJCAI** (pp. 209-226). London: William Kaufmann.

[Sloman 1978 1]
    Sloman, A. Sloman, A. (1978a). The Computer Revolution in Philosophy. Hassocks, Sussex: Harvester Press (and Humanities Press) Revised 2015. http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp

[Sloman 1978b]
    Sloman, A. (1978b). What About Their Internal Languages? Commentary on three articles in BBS Journal 1978, 1 (4). BBS , 1 (4), 515.

[Sloman 1979]
    Sloman, A. (1979). The primacy of non-communicative language. In M. MacCafferty & K. Gray (Eds.), The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979 (pp. 1-15). London: Aslib.

[Sloman 2005]
   Sloman, A. Sloman, A. (2005, September). Discussion note on the polyflap domain (to be explored by an
   'altricial' robot) (Research Note No. COSY-DP-0504). Birmingham, UK: School of Computer Science, University
   of Birmingham. Available from
   http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps

[Sloman 2013]
   Sloman, A. (2013). Meta-Morphogenesis and Toddler Theorems: Case Studies. School of Computer Science,
   The University of Birmingham. (Online discussion note) Available from http://goo.gl/QgZU1g

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

[Sloman Chappell 2007]
   Sloman, A. Chappell, J. (2007). Computational Cognitive Epigenetics (Commentary on (Jablonka & Lamb,
   2005)). BBS , 30 (4), 375-6.

[Tarsitano 2006]
   Tarsitano, M. (2006, December). Route selection by a jumping spider (Portia labiata) during the locomotory
   phase of a detour. Animal Behaviour , 72, Issue 6 , 1437-1442.

[Vetter 2013]
   Vetter, B. Vetter, B. (2013, Aug). 'Can' without possible worlds: semantics for anti-Humeans. Imprint
   Philosophers, 13 . (16)

[Weir, Chappell, Kacelnik . 2002]
   Weir, A A S., Chappell, J. Kacelnik, A. (2002). Shaping of hooks in New Caledonian crows. Science, 297 (9
   August 2002), 981.

[Whitehead Russell 1910-1913]
   Whitehead, A N. Russell, B. (1910-1913). Principia Mathematica Vols I - III. Cambridge: Cambridge University
   Press.

Footnotes:

[1] This is a snapshot of part of the Turing-inspired Meta-Morphogenesis project.

[2] I did not notice this "Polyflap stability theorem" until I tried to think of an example. I did not need to
do any experiments and collect statistics to recognize its truth (given familiar facts about gravity).
Do you?

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

[4] This video gives some details: https://www.youtube.com/watch?v=pjtioIFuNf8

[5]
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/chewing-test.html

[6] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision/plants
presents a botanical challenge for vision researchers.

[7] There seems to be uncertainty about dates and who contributed what. I'll treat Euclid as a figurehead for a tradition that includes many others, especially Thales, Pythagoras and Archimedes - perhaps the greatest of them all, and a mathematical precursor of Leibniz and Newton. More names are listed here:
https://en.wikipedia.org/wiki/Chronology_of_ancient_Greek_mathematicians I don't know much about mathematicians on other continents at that time or earlier. I'll take Euclid to stand for all of them, because of the book that bears his name.

[8] Moreover, it does not propagate misleading falsehoods, condone oppression of women or non-believers, or promote dreadful mind-binding in children.

[9] http://web.mnstate.edu/peil/geometry/C2EuclidNonEuclid/8euclidnoneuclid.htm

[10] My 1962 DPhil thesis [Sloman 1962] presented Kant's ideas, before I had heard about AI.
http://www.cs.bham.ac.uk/research/projects/cogaff/thesis/new

[11] http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html

[12] I was unaware of this until I found the Wikipedia article in 2015:
https://en.wikipedia.org/wiki/Angle_trisection#With_a_marked_ruler

[13] Much empirical research on number competences grossly over simplifies what needs to be explained, omitting the role of reasoning about 1-1 correspondences.

[14] Richard Gregory demonstrated that a 3-D structure can be built that looks exactly like an impossible object, but only from a particular viewpoint, or line of sight.

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