

**Invited talk at International Workshop on
Artificial Intelligence and Cognition (AIC 2014)**

<http://aic2014.di.unito.it/>

<http://aic2014.di.unito.it/program.html>

November 26-27 University of Turin, Italy

**How can we reduce the gulf between
artificial and natural intelligence?**

First characterise the gulf accurately!

Aaron Sloman

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

University of Birmingham, UK

NOTE:

A revised version of the extended abstract will be published
in the workshop proceedings after the conference. Comments, criticisms and suggestions for
improvement are all welcome. (a.sloman [AT] cs.bham.ac.uk)

Abstract

There are growing numbers of impressive successes of artificial intelligence and robotics, many of
them summarised at <http://aitopics.org/news>.

Yet there remain huge chasms between artificial systems and forms of natural intelligence in humans
and other animals -- including weaver-birds, elephants, squirrels, dolphins, orangutans, carnivorous
mammals, and their prey.

(Sample weaver bird cognition here: <http://www.youtube.com/watch?v=6svAIgEnFvw>.)

Fashionable "paradigms" offering definitive answers come and go (sometimes reappearing with new
labels). Yet no AI or robotic systems come close to modelling or replicating the development from
helpless infant over a decade or two to plumber, cook, trapeze artist, bricklayer, seamstress, dairy
farmer, shop-keeper, child-minder, professor of philosophy, concert pianist, mathematics teacher,
quantum physicist, waiter in a busy restaurant, etc. Human and animal developmental trajectories
vastly outstrip, in depth and breadth of achievement, the products of artificial learning systems,
although AI systems sometimes produce super-human competences in restricted domains, such as
proving logical theorems, winning at chess or Jeopardy, and perhaps playing table tennis at
championship level one day in the distant future? (<http://www.youtube.com/watch?v=tIIJME8-au8>).

I'll outline a very long-term multi-disciplinary research programme addressing these and other
inadequacies in current AI, robotics, psychology, neuroscience and philosophy of mathematics and
mind, in part by building on past work, and in part by looking for very different clues and challenges:
the Meta-Morphogenesis project, partly inspired by Turing's work on morphogenesis.

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

Note:

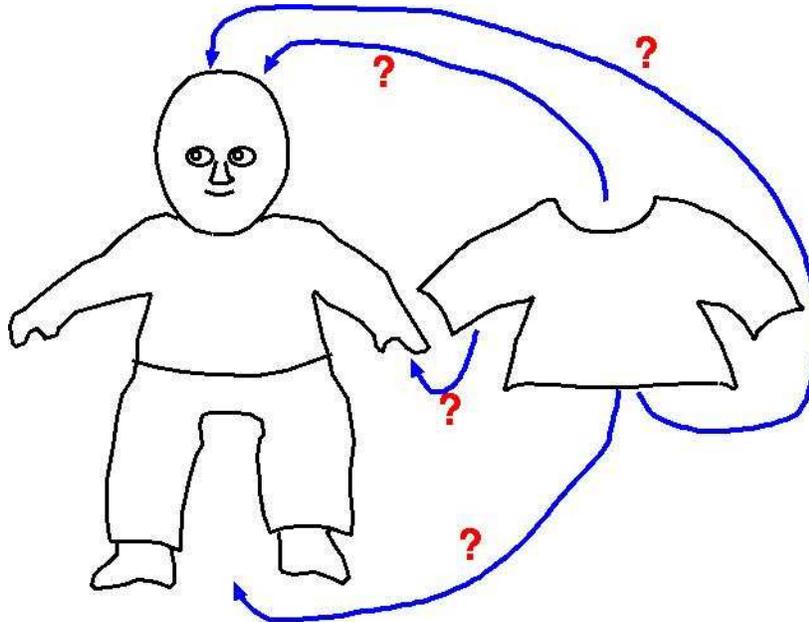
This project is unfunded and I have no plans to apply for funding, though others may do so if they wish.

Extended Abstract

WORK IN PROGRESS -- LIABLE TO CHANGE

(Installed 20 Sep 2014. Last updated: 30 Sep 2014)

There are many ways in which current robots and AI systems fall short of the intelligence of humans and other animals, including their ability to reason about topology and continuous deformation (for examples see [Sauvy and Sauvy 1974](#) and [this document](#)). Don't expect any robot (even with soft hands and compliant joints) to be able to dress a two year old child (safely) in the near future, a task that requires understanding of both topology and deformable materials, among other things. (As illustrated in [this video](#).) Understanding why things work and don't work lags even further behind abilities to perform tasks, often achieved by programming or training. For example, understanding why it's not a good idea to start putting on a shirt by inserting a hand into a cuff and pulling the sleeve up over the arm requires a combination of topological and metrical reasoning: -- a type of mathematical child-minding theorem, not taught in schools but understood by most child-minders, even if they have never articulated the theorem and cannot articulate the reasons why it is true. Can you? Merely pointing at past evidence showing that attempts to dress a child that way always fails does not explain why it is impossible.



What sequence of movements could get the shirt onto the child if the shirt is made of material that is flexible but does not stretch much? Why would it be a mistake to start by pulling the cuff over the hand, or pushing the head through the neck-hole? What difference would it make if the material could be stretched arbitrarily without being permanently changed?

In more obviously mathematical domains, where computers are commonly assumed to excel, the achievements are narrowly focused on branches of mathematics using inference methods based on arithmetic, algebra, logic, probability and statistical theory.

The "generative grammar" of spatial structures and processes is rich and deep, and is not concerned only with linear sequences discrete sequences. In fact there are multiple overlapping space-time grammars, involving different collections of objects assembled, disassembled, moved, repaired, etc. and used, often for many purposes and in many ways.

There are different overlapping subsets of spatio-temporal possibilities, with different mathematical structures, including Euclidean and non-Euclidean geometries (e.g. the geometry of the surface of a hand, or face is non-euclidean) and various subsets of topology. Mechanisms for for acquiring and using these "possibility subsets", i.e. possible action sequences and trajectories seem to be used by pre-verbal children and other animals. That suggests that those abilities, must have evolved before linguistic capabilities. They seem to be at work in young children playing with toys before they can understand or speak a human language. The starting capabilities extended through much spatial exploration, seem to provide much of the subject matter (semantic content) for many linguistic communications.

The early forms of reasoning and learning in young humans, and corresponding subsets in other animals, are beyond the scope of current AI theorem provers, planners, reasoners, or learning systems that I know of. Those forms seem to be used by non-human intelligent animals that are able to perceive both possibilities and constraints on possibilities in spatial configurations. Betty, a New Caledonian crow, made headline news in 2002 when she surprised Oxford researchers by making a hook from a straight piece of wire, in order to lift a bucket of food out of a vertical glass tube. Moreover, in a series of repeated challenges she made multiple hooks, using at least four very different strategies, taking advantage of different parts of the environment, all apparently in full knowledge of what she was doing and why -- as there was no evidence of random trial and error behaviour. Why did she not go on using the earlier methods, which all worked? Several of the videos showing the diversity of techniques are still available here: <http://users.ox.ac.uk/~kgroup/tools/movies.shtml>

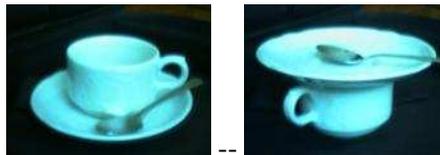
It is very unlikely that you have previously encountered and solved the problem posed below the following image, yet many people very quickly think of a solution.



Suppose you wanted to use one hand to lift the mug to a nearby table without any part of your skin coming into contact with the mug, and without moving the book on which the mug is resting, what could you do, using only one hand?

In order to think of a strategy you do not need to know the exact, or even the approximate, sizes of the objects in the scene, how far away they are from you, exactly what force will be required to lift the mug, and so on. It may occur to you that if the mug is full of liquid and you don't want to spill any of it, then a quite different solution is required. (Why? Is there a solution?).

Another set of example action strategies is provided by the following two images.



Consider one or more sequences of actions that would enable you to change the physical configuration depicted on the left into the configuration depicted on the right -- not necessarily in exactly the same locations as the objects

depicted. Then do the same for the actions required to transform the right configuration to the left one.

At how many different levels of abstraction can you think of the process, where the levels differ in the amount of detail (e.g. metrical detail) of each intermediate stage. For example, when you first thought about the problem did you specify which hands or which fingers would be used at every stage? If you specified the locations used to grasp the cup, the saucer and the spoon, what else would have to change to permit those grasps. The point about all this is that although you do not normally think of using mathematics for tasks like this, if you choose a location at which to grasp the cup using finger and thumb of your left hand, that will constrain the 3-D orientation of the gap between between finger and thumb, if you don't want the cup to be rotated by the fact of bringing finger and thumb together.

When you think about such things even with fairly detailed constraints on the possible motions, you will not be thinking about either the nervous signals sent to the muscles involved, nor the patterns of retinal stimulation that will be provided -- and in fact the same actions can produce different retinal processes depending on the precise position of the head, and the direction of gaze of the eyes, and whether and how the fixation changes during the process. Probably the fixation requirements will be more constrained for a novice at this task than for an expert, but either way humans, and I suspect other animals, do not need to reason about such details if they use an ontology of 3-D structures and processes rather than an ontology of sensory and motor brain signals. Contrast this with the sorts of assumptions discussed in [Clark \(2013\)](#), and many others who attempt to build theories of cognition on the basis of sensory-motor control loops.

This ability to think about sequences of possible alterations in a physical configuration without actually doing anything, and without having full metrical information, inspired much early work in AI, including the sorts of symbolic planning used by Shakey the Stanford robot and Freddy the Edinburgh robot, though at the time the technology available (including available computer power) was grossly inadequate for the task, including ruling out visual servo control of actions.

The capabilities and information-processing requirements supporting them that I have been discussing are ignored by researchers who claim that intelligent robots require only the right physical mode of interaction with the environment, not central reasoning capabilities. Some have supported this claim by referring to or demonstrating "passive walkers" <http://www.youtube.com/watch?v=N64KOQkbyiI>. Try putting a brick in front of a passive walker. I'll refrain from naming and embarrassing authors who, for a while, enthusiastically used such toys as pointers to a "New artificial intelligence", using labels such as "embodied", "enactivist", "behaviour based", and "situated", to characterise their new acclaimed paradigm -- not realising that those approaches are at least as selective as the older reasoning based approaches that they criticised. (Some of that history was presented in [Boden \(2006\)](#)).

The requirements for perception and action mechanisms differ according to which "central" layers the organism has. For instance, for an organism able to use deliberative capabilities to think of evaluate and select multi-step plans, where most of the actions will occur in situations that do not exist yet, it is not enough to identify objects and their relationships (pencil, mug, handle of mug, book, window-frame, etc.) in a current visual percept. It is also necessary to be able to "think ahead" about possible actions at a suitable level of abstraction, derived from what is perceived.

This ability to reason about possible actions at a level of generality that abstracts from metrical details seems to be closely related to the abilities of ancient Greeks to make mathematical discoveries about possible configurations of lines and circles and the consequences of changing those configurations, without being tied to particular lengths, angles, curvatures, etc., in Euclidean geometry or topology. As far as I know, no current robot can do this, and neuroscientists don't know how brains do it. Some

examples of mathematical reasoning that could be related to reasoning about practical tasks and which are currently beyond what AI reasoners can do, are presented here:

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

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

In 1971 I presented a paper at an AI conference (IJCAI) arguing that the focus solely on logic-based reasoning could hold up progress in AI, because it ignored forms of spatial reasoning that had proved powerful in mathematics and practical problem solving. I did not realise then how difficult it would be to explain exactly what the alternatives were and how they worked -- despite many conferences and journal papers on diagrammatic reasoning since then. There have also been several changes of fashion promoted by various AI researchers (or their critics) including use of neural nets, constraint nets, evolutionary algorithms, dynamical systems, behaviour-based systems, embodied cognition, situated cognition, enactive cognition, autopoiesis, morphological computation, statistical learning, bayesian nets, and probably others that I have not encountered, often accompanied by hand-waving and hyperbole without much science. In parallel with this there has been continued research advancing older paradigms for symbolic and logic based, theorem proving, planning, and grammar based language processing. Several of the debates are analysed in Boden(2006).[REF]

There are many other inadequacies in current AI, including, for example the lack of an agreed framework for relating information-processing architectures to requirements in engineering contexts or to explanatory models in scientific contexts (e.g. attempts to model emotions, or learning capabilities, in humans or other animals. I also think we are using a seriously restricted set of forms of representation (means of encoding information) partly because of the educational backgrounds of researchers (as a result of which many of them assume that spatial structures must be represented using mechanisms based on Cartesian coordinates) and partly because of a failure to analyse in sufficient detail the problems overcome by many animals in their natural environments.

Standard research techniques are not applicable to the study of such capabilities in young children and other animals because there is so much individual variation, but the widespread availability of cheap video cameras has led to a large and growing collection of freely available examples.

However, researchers have to learn what to look for. For example, *online* intelligence requires highly trained precisely controlled responses matched to fine details of the physical environment, e.g. catching a ball, playing table tennis, picking up a box and putting it on another. In contrast *offline* intelligence involves understanding not just existing spatial configurations but also the possibilities for change and constraints on change, and for some tasks the ability to find sequences of possible changes to achieve a goal, where some of the possibilities are not specified in metrical detail because they do not yet exist. This requires the ability to construct relatively abstract forms of representation of perceived or remembered situations to allow plans to be constructed with missing details that can be acquired later during execution. You can think about making a train trip to another town without having information about where you will stand when purchasing your ticket or which coach you will board when the train arrives. You can think about how to rotate a chair to get it through a doorway without needing information about the precise 3-D coordinates of parts of the chair or knowing exactly where you will grasp it, or how much force you will need to apply at various stages of the move.

There is no reason to believe that humans and other animals have to use probability distributions over possible precise metrical values. Even thinking about such precise values probabilistically is highly unintelligent when reasoning about topological relationships or partial orderings (nearer, thinner, a bigger angle, etc.) is all that's needed, as I have tried to illustrate here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html> Unfortunately, the mathematically sophisticated, but nevertheless unintelligent, modes of thinking are used in many

robots, after much statistical learning (to acquire probability distributions) and complex probabilistic reasoning, that is potentially explosive. That is in part a consequence of the unjustified assumption that all spatial properties and relations have to be expressed in Cartesian coordinate systems. Human mathematicians did not know about them when they proved their first theorems about Euclidean geometry.

NOTE to be updated:

It is clear that the earliest spatial cognition could not have used full euclidean geometry, including its uniform metric. I suspect that the metrical version of geometry was a result of a collection of transitions adding richer and richer non-metrical relationships, including networks of partial orderings of size, distance, angle, speed, curvature, etc. Later indefinitely extendable partial metrics were added: distance between X and Y is at least three times the distance between P and Q and at most five times that distance. Such procedures could allow previously used standards to sub-divided with arbitrarily increasing precision. At first this must have been applied only to special cases, then later somehow (using what cognitive mechanisms) extrapolated indefinitely, implicitly using a Kantian form of **potential** infinity (long before Kant realised the need for this). Filling in the details of such a story, and relating it to varieties of cognition not only in the ancestors of humans but also many other existing species will be a long term multi-disciplinary collaborative task, with deep implications for neuroscience, robotics, psychology, philosophy of mathematics and philosophy of mind. (Among others.)

Moreover, human toddlers appear to be capable of making proto-mathematical discoveries ("toddler theorems") even if they are unaware of what they have done. But the process starts in infancy, e.g. when I saw a 11 month old infant discover, apparently with great delight, that she could hold a ball between her upturned foot and the palm of her hand. (Picture to be added.)

Animal abilities to perceive and use complex novel affordances appear to be closely related. Not only computational models, but also current psychology and neuroscience, don't seem to come close -- especially if we consider not only simple numerical mathematics, on which many psychological studies of mathematics seem to focus, but also topological and geometrical reasoning, and the essentially mathematical ability to discover a generative grammar closely related to the verbal patterns a child has experienced in her locality, where the grammar is very different from those discovered by children exposed to thousands of other languages.

There seem to be key features of some of those developmental trajectories that could provide clues, including some noticed by Piaget and his former colleague, Annette Karmiloff-Smith. (EXPAND)

Partly inspired by one of Alan Turing's last papers (The Chemical Basis of Morphogenesis) (1952), I'll outline a very long term collaborative project for building up an agreed collection of explanatory tasks, and present some ideas about what has been missed in most proposed explanatory theories. Perhaps researchers who disagree, often fruitlessly, about what the answers are can collaborate fruitfully on finding out what the questions are, sine what needs to be explained is far from obvious. The Meta-Morphogenesis project is concerned with trying to understand what varieties of information processing biological evolution has achieved, not only in humans but across the spectrum of life. Many of the achievements are far from obvious.

A more detailed, but still evolving, introduction to the project can be found here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

REFERENCES

(To be extended.)

- [Can a robot grasp grasping?](#) (This web site.)
How can a robot understand what's going on when it grasps something?
- M. A. Boden, 2006
Mind As Machine: A history of Cognitive Science (Vols 1--2) (2006) O.U.P
- A. M. Turing, 1952, The Chemical Basis Of Morphogenesis,
Phil. Trans. R. Soc. London B 237, 237, pp. 37--72,
- Andy Clark, 2013,
Whatever next? Predictive brains, situated agents, and the future of cognitive science, *Behavioral and Brain Sciences*, 36, 3, pp. 1--24,
- Jean Sauvy and Simonne Sauvy, with an introduction by Bill Brookes
The Child's Discovery of Space: From hopscotch to mazes -- an introduction to intuitive topology,
Penguin Education, Harmondsworth, 1974. Translated from the French by Pam Wells,
-
-

Maintained by [Aaron Sloman](#)
[School of Computer Science](#)
[The University of Birmingham](#)