NOTES FOR
REMOTE BICA TUTORIAL, MOSCOW 3rd Aug 2017
https://goo.gl/NWyC35

Expanded version of notes for:
AISB17 Symposium on Computing and Philosophy
Bath University 20th April 2017

Gaps Between Human and Artificial Mathematics

Some deep, largely unnoticed, gaps in current AI, and what Alan Turing might have done about them.

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

Alternative title:
THE SELF-INFORMING UNIVERSE

"Systems with self-improving theories that automatically increase their understanding of the world around them."
Jiri Wiedermann (AISB 2017): See
http://aisb2017.cs.bath.ac.uk/conference-edition-proceedings.pdf

---

NOTE:
A revised extended version of this document, including a video recording, is available at:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html
(Invited, remotely presented, contribution to the IJCAI Workshop on Architectures for Generality and Autonomy, Melbourne Australia on 19 Aug 2017: http://cadia.ru.is/workshops/aga2017/)

(DRAFT: Liable to change)

Installed: 18 Apr 2017
Last updated: 11 Jul 2017 (for AISB); 3 Aug 2017 (for BICA); 7 Aug 2017; 29 Aug 2017
Available as html and pdf:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/aisb-CandP.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/aisb-CandP.pdf

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

---

Partial progress report on the Turing-inspired Meta-Morphogenesis project

> Trying to understand intelligence by studying only human intelligence
> is as misguided as trying to understand life by studying only human life.

Background: What is mathematical discovery? (Euclid, Kant and Einstein)
This work started before I heard about Artificial Intelligence or learnt to program. After a degree in mathematics and physics at Cape Town, I came to Oxford in October 1957, intending to do research in mathematics (after further general study). Because I did not like some of the compulsory mathematics courses (e.g. fluid dynamics) I transferred from mathematics to Logic with Hao Wang as my supervisor, and and became friendly with philosophy graduate students, with whom I used to argue. This eventually caused me to transfer to Philosophy. I am still trying to answer the questions about mathematical knowledge that drove me at that time.

The philosophers I met (mostly philosophy research students) were mistaken about the nature of mathematical discovery as I had experienced it while doing mathematics. E.g. some of them accepted David Hume's categorisation of claims to knowledge, which seemed to me to ignore important aspects of mathematical discovery.

---

1. Hume's first category was "abstract reasoning concerning quantity or number", also expressed as knowledge "discoverable by the mere operation of thought". This was sometimes thought to include all "trivial knowledge" consisting only of relations between our ideas, for example, "All bachelors are unmarried". Kant labelled this category of knowledge "Analytic".

   It is sometimes specified as knowledge that can be obtained by starting from definitions of words and then using only pure logical reasoning, e.g. *"No bachelor uncle is an only child"*.

2. Hume's second category was empirical knowledge gained, and tested, by making observations and measurements i.e. "experimental reasoning concerning matter of fact and existence". This would include much common sense knowledge, scientific knowledge, historical knowledge, etc.

3. His third category was everything that could not fit into either the first or second. He described the residue as "nothing but sophistry and illusion" urging that all documents claiming such knowledge should be "committed to flames". I assume he was thinking mainly of metaphysics and theology.

Warning: I am not a Hume scholar. For more accurate and more detailed summaries of his ideas search online. e.g.
https://en.wikipedia.org/wiki/David_Hume
https://plato.stanford.edu/entries/hume/

The philosophers I met seemed to believe that all mathematical knowledge was in Hume's first category and was therefore essentially trivial. (My memory is a bit vague about 60 year old details.)

But I knew from my own experience of doing mathematics that mathematical knowledge did not fit into any of these categories: it was closest to the first category, but was not trivial, and did not come only from logical deductions from definitions.

I then discovered that Immanuel Kant had criticised Hume for not allowing a category of knowledge that more accurately characterised mathematical knowledge, in his 1781 book, *"Critique of Pure Reason"*.

But the philosophers thought Kant's ideas about mathematical knowledge being non-trivial and non-empirical were mistaken because he took knowledge of Euclidean geometry as an example. They thought Kant had been proved wrong when Einstein and Eddington showed that space was not Euclidean, by demonstrating the curvature of light rays passing close to the sun:
https://en.wikipedia.org/wiki/Euclidean_geometry#20th_century_and_general_relativity

This argument against Kant was misguided for several reasons. In particular it merely showed that human mathematicians could make mistakes, e.g. by thinking that 2D and 3D spaces were necessarily Euclidean.

> In a Euclidean plane surface, if **P** is any point, and **L** any straight line that does not pass through **P**,
> there will be exactly one straight line through **P** in the plane, that never intersects **L**.
> I.e. there is a unique line through **P** and parallel to **L**.

I don't think anything Kant wrote implied that mathematicians are infallible. The extent of their fallibility was illustrated by Lakatos in his *Proofs and Refutations* (1976))

Moreover, before Einstein's work, mathematicians had previously discovered that not all spaces are necessarily Euclidean and that there were different kinds of space in which the parallel axiom was false (elliptical and hyperbolic spaces). If Kant had known this, I am sure he would have changed the examples that assumed the parallel axiom. Removing it leaves enough rich and deep mathematical content to illustrate Kant's claims, including the mathematical discovery that a Euclidean geometry without the parallel axiom is consistent with both Euclidean and non-Euclidean spaces: as good an example of a non-analytic necessary truth as any Kant presented.

He could have used the discovery that Euclidean geometry without the parallel axiom could be extended in three different ways with very different consequences as one of his examples of a mathematical discovery that is not derivable from definitions by logic, and is a necessary truth, and can be discovered by mathematical thinking, and does not need empirical tests at different locations, altitudes, or on different planets, etc.

In 1962 I completed my DPhil thesis defending Kant, now online Sloman(1962)

I went on to become a lecturer in philosophy, but I was left feeling that my thesis did not answer all the questions, and something more needed to be done. So when Max Clowes, a pioneering AI vision researcher came to Sussex university and introduced me to AI and programming I was eventually persuaded to try to show how AI could support Kant, by demonstrating how to build a "baby robot" that "grows up" to make new mathematical discoveries in roughly the manner that Kant had described, including replicating some of the discoveries of ancient mathematicians like Archimedes, Euclid and Pythagoras.

Max Clowes died in 1981. A tribute to him with annotated bibliography is here. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61

This would require a form of learning from totally different from both

- the methods based on exploring logical consequences of axioms as done in AI theorem provers
- the methods based on collecting statistical evidence and performing probabilistic reasoning, as done in systems that use "deep learning" and and forms of probabilistic reasoning.

The latter methods are logically incapable of demonstrating truths of mathematics, which are concerned with necessities and impossibilities, not mere probabilities.

(Including some that human toddlers and intelligent non-human species seem able to discover, even if unwittingly, as I have tried to demonstrate, e.g. in this partial survey

of what I now call "toddler theorems":
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html)

Part of my argument in the thesis, inspired by Kant, was that intelligent robots, like intelligent humans, needed forms of mathematical reasoning that were not restricted to use of logical derivations from definitions, and were also different from empirical reasoning based on experiment and observation.

Encouraged by Max Clowes I published a paper (at IJCAI 1981) that challenged the "logicist" approach to AI proposed by John McCarthy, one of the founders of AI, as presented in McCarthy and Hayes (1969). My critique of logicism, emphasising the heuristic benefits of "analogical" representations is Sloman (1971).

As a result I was invited to spend a year (1972-3) doing research in AI at Edinburgh University. I hoped it would be possible to use AI to defend Kant's philosophical position by showing how to build a "baby robot" without mathematical knowledge, that could grow up to be a mathematician in the same way as human mathematicians did, including, presumably the great ancient mathematicians who knew nothing about modern logic, formal systems of reasoning based on axioms (like Peano's axioms for arithmetic) and did not assume that geometry could be modelled in arithmetic as Descartes had shown.

I published a sort of "manifesto" about this in 1978 (*The Computer Revolution in Philosophy*, freely available online, with additional notes and comments.)

The task turned out to be much more difficult than I had expected and now nearly 40 years later, after doing a lot of work in AI, including a lot of work on architectures for intelligent agents,

http://www.cs.bham.ac.uk/research/projects/cogaff/

a toolkit for exploring alternative agent architectures,

http://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html

work on requirements for human-like vision systems, and many related topics, I am still puzzled about exactly what is missing from AI.

Since 2012, as explained later, I have been trying to fill the gaps by means of the Turing-inspired Meta-Morphogenesis project, a very difficult long term project, which I suspect Alan Turing was thinking about in the years before he died, in 1954.

In parallel with this I am trying to analyse the forms of reasoning required for the ancient mathematical discoveries in geometry and topology (illustrated below), with the aim eventually of specifying detailed requirements for a machine to make such discoveries. That may give new clues regarding how animal brains work.

The Meta-Morphogenesis project

The Turing-inspired Meta-Morphogenesis project was proposed in the final commentary in *Alan Turing - His Work and Impact*, a collection of papers by and about Turing published on the occasion of his centenary[6].

The project defines a way of trying to fill gaps in our knowledge concerning evolution of biological information processing that may give clues regarding forms of computation in animal brains that have not yet been re-invented by AI researchers.

This may account for some of the enormous gaps between current AI and animal intelligence, including gaps between mathematical abilities of current AI systems and the abilities of ancient mathematicians whose discoveries are still being used all over world, e.g. Archimedes, Euclid, Pythagoras and Zeno.

Evolution of information processing capabilities and mechanisms is much harder to study than evolution of physical forms and physical behaviours, e.g. because fossil records can provide only very indirect evidence regarding information processing in ancient organisms. Moreover it is very hard to study all the internal details of information processing in current organisms. Some of the reasons will be familiar to programmers who have struggled to develop debugging aids for very complex multi-component AI virtual machines.

Because we cannot expect to find fossil records of information processing, or the mechanisms used, the work has to be highly speculative. But conjectures should be constrained where possible by things that are known. Ideally these conjectures will provoke new research on evolutionary evidence and evidence in living species. However, as often happens in science, the evidence may not be accessible with current tools. Compare research in fundamental physics (e.g. Tegmark (2014)).

The project presents challenges both for the theory of biological evolution by natural selection, and for AI researchers aiming to replicate natural intelligence, including mathematical intelligence. This is a partial progress report on a long term attempt to meet the challenges. A major portion of the investigation at this stage involves (informed) speculation about evolution of biological information processing, and the mechanisms required for such evolution, including evolved construction-kits. The need for which has not been widely acknowledged by evolutionary theorists.

An extended abstract for a closely related invited talk at the AISB Symposium on computational modelling of emotions is also available online at:
http://www.cs.bham.ac.uk/research/projects/cogaff/aisb17-emotions-sloman.pdf

---

The Meta-Morphogenesis Project

This is a partial progress report on the Meta-Morphogenesis (M-M) Project -- also called The Self-Informing Universe project, originally proposed during the Turing Centenary year (2012).

A lot of work has been done on the project since then, some of it summarised below, especially the developing theory of evolved construction kits of various sorts Sloman[2017], but there are still many unsolved problems, both about the processes of evolution and the products in brains of intelligent animals.

I am not primarily interested in AI as engineering: making useful new machines. Rather I want to understand how animal brains work, especially animals able to make mathematical discoveries like the amazing discoveries reported in Euclid's *Elements* over 2000 years ago.

My interest in AI (which started around 1969) and my work on the The M-M project, originally came from my interest in defending Immanuel Kant's philosophy of mathematics in his (1781), and partly on a conjectured answer to the question: 'What would Alan Turing have worked on if he had not died two years after publication of his 1952 paper on Chemistry and Morphogenesis (Turing 1952). This is now the most cited of his publications. though largely ignored by philosophers, cognitive scientists and AI researchers.

I suspect that if Turing had lived several decades longer, he would have tried to understand forms of information processing needed to control behaviour of increasingly complex organisms produced by evolution, starting from the very simplest forms produced somehow on a lifeless planet produced by condensed gaseous matter and dust particles. That is the M-M project.

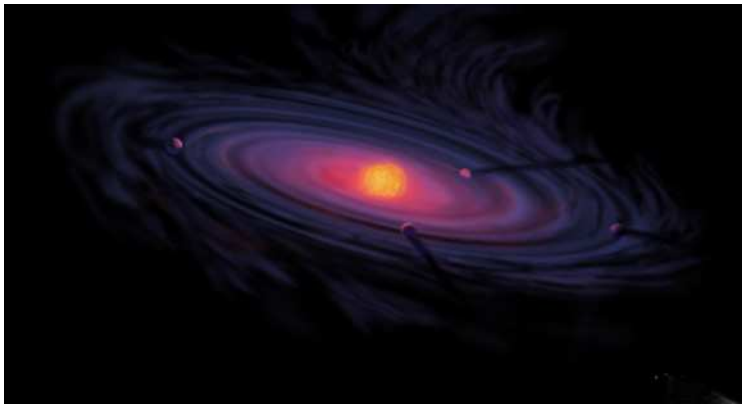

[NASA artist's impression of a protoplanetary disk, from WikiMedia]

How could this come about?

I have nothing to add to conjectures by others about the initial, minimal forms of life, e.g. see Ganti (2003).

However, controlled production of complex behaving structures needs increasingly sophisticated information processing:
-- in processes of reproduction, growth and development
-- for control of behaviour of complex organisms reacting to their environment, including other organisms.
(Regarding mechanisms for storing information required for reproduction Schrödinger (1944) had some profound observations.)

In simple organisms, control mainly uses presence or absence of sensed matter to turn things on or off or sensed scalar values to specify and modify other values (e.g. chemotaxis).

But as organisms and their internal structures become more complex, the need for structural rather than metrical specifications increases.

> Many artificial control systems are specified using collections of differential equations relating such measures. One of several influential attempts to generalise these ideas is the 'Perceptual Control Theory (PCT)' of William T Powers.

But use of numerical/scalar information is not general enough: It doesn't suffice for linguistic (e.g. grammatical or semantic) structures or for reasoning about topological relationships, or processes of structural change e.g. in chemical reactions or engineering assembly processes -- including 'toy' engineering, such as playing with meccano sets. It also cannot describe growth of organisms, such as plants and animals, in which new materials, new substructures, new relationships and new capabilities form -- including new information processing capabilities.

> For example, the changes between an egg and a chicken cannot be described by changes in a state-vector. Why not?

Turing's Morphogenesis paper [31] also focused on mechanisms (e.g. diffusion of chemicals) representable by scalar (numerical) changes, but the results included changes of structure described in words and pictures. As a mathematician, a logician and a pioneer of modern computer science he was well aware that the space of information-using control mechanisms is not restricted to numerical control systems.

> For example a Turing machine's operation involves changing linear sequences of distinct structures, not numerical measures.

In the last half century human engineers have discovered, designed and built additional increasingly complex and varied forms of control in interacting physical and virtual machines.

That includes control based on

> grammars, parsers, planners, reasoners, rule interpreters, problem solvers and many forms of automated discovery and learning.

Long before that, biological evolution produced and used increasingly complex and varied forms of information in construction, modification and control of increasingly complex and varied behaving mechanisms.

CONJECTURE:

> If Turing had lived several decades longer, he might have produced new theories about many intermediate forms of information in living systems and intermediate mechanisms for information-processing: intermediate between the very simplest forms and the most sophisticated current forms of life.

This would fill gaps in standard versions of the theory of natural selection. E.g. , the theory does not explain what makes possible the many forms of life on this planet, and all the mechanisms they use, including the forms that might have evolved in the past or may evolve in the future.

It merely assumes such possibilities and explains how a subset of realised possibilities persist and consequences that follow.

For example, the noted biologist Graham Bell wrote in ’Living complexity cannot be explained except through selection and does not require any other category of explanation whatsoever’Bell(2008).

Only a few defenders of Darwinian evolution seem to have noticed the need to explain

(a) what mechanisms make possible all the options between which choices are made, and

(b) how what is possible changes, and depends on previously realised possibilities.

CONJECTURE: USES OF EVOLVED CONSTRUCTION KITS

A possible defence of Darwinian evolution would enrich it to include investigation of

> (a) the Fundamental Construction Kit (FCK) provided by physics and chemistry before life existed,

(b) the many and varied ’Derived construction kits’ (DCKs) produced by combinations of natural selection and other processes, including asteroid impacts, tides, changing seasons, volcanic eruptions and plate tectonics.

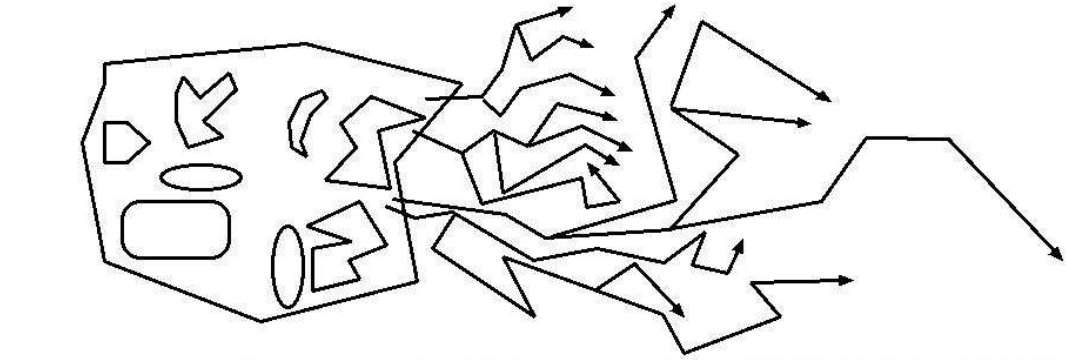

The Fundamental Construction Kit (FCK)

Evolutionary and other trajectories
from the FCK through the space of possibilities

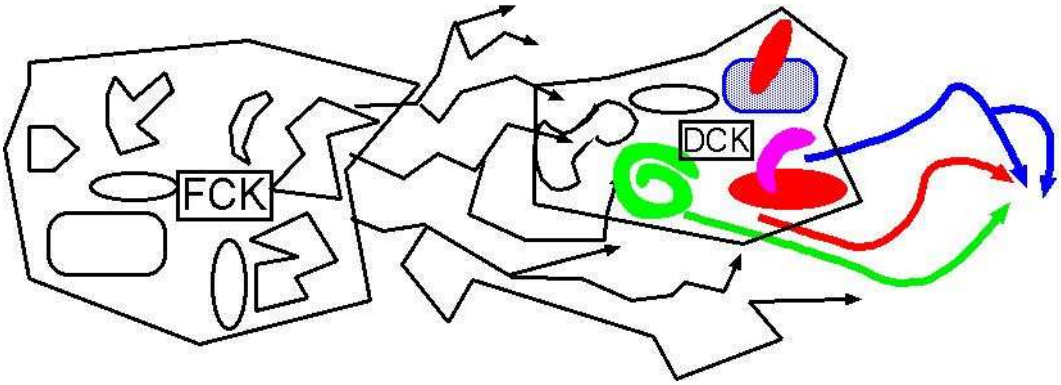

As new, more complicated, life forms evolved, with increasingly complex bodies, increasingly complex changing needs, increasingly broad behavioural repertoires, and richer branching possible actions and futures to consider, their information processing needs and opportunities also became more complex.

Somehow the available construction kits also diversified, in ways that allowed

construction not only of new biological materials and body mechanisms, supporting new more complex and varied behaviours

but also

new more sophisticated information-processing mechanisms, enabling organisms, either alone or in collaboration, to deal with increasingly complex challenges and opportunities.

DEEP DESIGN DISCOVERIES

Many deep discoveries were made by evolution, including designs for DCKs that make possible new forms of information processing.

These have important roles in animal intelligence, including perception, conceptual development, motivation, planning, and problem solving, including

-- topological reasoning about properties of geometrical shapes and shape-changes.
-- reasoning about possible continuous rearrangements of material objects (much harder than planning moves in a discrete space).

Different species, with different needs, habitats and behaviours, use information about different topological and geometrical relationships, including

-- birds that build different sorts of nests,
-- carnivores that tear open their prey in order to feed,
-- human toddlers playing with (or sucking) body-parts, toys, etc.

Later on, in a smaller subset of species (perhaps only one species?) new meta-cognitive abilities gradually allowed previous discoveries to be noticed, reflected on, communicated, challenged, defended and deployed in new contexts.

Such 'argumentative' interactions may have been important precursors for chains of reasoning, including the proofs in Euclid's Elements.

WHY IS THIS IMPORTANT?

This is part of an attempt to explain how it became possible for evolution to produce mathematical reasoners.

New deep theories, explanations, and working models should emerge from investigation of preconditions, biological and technological consequences, limitations, variations, and supporting mechanisms for biological construction kits of many kinds.

For example, biologists have pointed out that specialised construction kits, sometimes called 'toolkits', supporting plant development were produced by evolution, making upright plants possible on land (some of which were later found useful for many purposes by humans, e.g. ship-builders).

Specialised construction kits were also needed by vertebrates and others by various classes of invertebrate forms of life.

INFORMATION PROCESSING

Construction kits for biological information processing have received less attention.

One of the early exceptions was Schrödinger's little 1944 book *What is life?*

More general construction kits that are tailorable with extra information for new applications can arise from discoveries of parametrisable sub-spaces in the space of possible mechanisms

e.g. common forms with different sizes, or different ratios of sizes, of body parts, different rates of growth of certain body parts, different shapes or sizes of feeding apparatus, different body coverings, etc.

Using a previously evolved construction kit with new parameters (specified either in the genome, or by some aspect of the environment during development) can produce new variants of organisms in a fraction of the time it would take to evolve that type from the earliest life forms.

Similar advantages have been claimed for the use of so-called Genetic Programming (GP) using evolved, structured, parametrised abstractions that can be re-deployed in different contexts, in contrast with Genetic Algorithms (GAs) that use randomly varied flat strings of bits or other basic units.

Evolution sometimes produces specifications for two or more different designs for different stages of the same organism, e.g. one that feeds for a while, and then produces a cocoon in which materials are transformed into a chemical soup from which a new very different adult form (e.g. butterfly, moth, or dragon fly) emerges, able to travel much greater distances than the larval form to find a mate or lay eggs.

These species use mathematical commonality at a much lower level (common molecular structures) than the structural and functional designs of larva and adult, in contrast with the majority of organisms, which retain a fixed, or gradually changing, structure while they grow after hatching or being born, but not fixed sizes or size-ratios of parts, forces required, etc.

Mathematical discoveries were implicit in evolved designs that support parametrisable variable functionalities, such as evolution's discovery of homeostatic control mechanisms that use negative feedback control, billions of years before the Watt centrifugal governor was used to control speed of steam engines.13 Of course, most instances of such designs would no more have any awareness of the mathematical principles being used than a Watt-governor, or a fan-tail windmill (with a small wind-driven wheel turning the big wheel to face the wind) does.

In both cases a part of the mechanism acquires information about something (e.g. whether speed is too high or too low, or the direction of maximum wind strength) while another part does most of the work, e.g. transporting energy obtained from heat or wind power to a new point of application.

Such transitions and decompositions in designs could lead to distinct portions of genetic material concerned with separate control functions, e.g. controlling individual development and controlling adult use of products of development, both encoded in genetic material shared across individuals.

METACOGNITION EVOLVES

Very much later, some meta-cognitive products of evolution allowed individuals (humans, or precursors) to attend to their own information-processing (essential for debugging), thereby 'rediscovering' the structures and processes, allowing them to be organised and communicated -- in what we now call mathematical theories, going back to Euclid and his predecessors (about whose achievements there are still many unanswered questions).

If all of this is correct then the physical universe, especially the quantum mechanical aspects of chemistry discussed by Schrödinger provided not only

> a construction kit for genetic material implicitly specifying design features of individual organisms,

but also

> a 'Fundamental' construction kit (FCK) that can produce a wide variety of 'derived' construction kits (DCKs)

some used in construction of individual organisms, others in construction of new, more complex DCKs, making new types of organism possible.

Moreover, as Schrödinger and others pointed out, construction kits that are essential for micro-organisms developing in one part of the planet can indirectly contribute to construction and maintenance processes in totally different organisms in other locations, via food chains, e.g. because most species cannot synthesise the complex chemicals they need directly from freely available atoms or subatomic materials. So effects of DCKs can be very indirect.

Functional relationships between the smallest life forms and the largest will be composed of many sub-relations.

Such dependency relations apply not only to mechanisms for construction and empowerment of major physical parts of organisms, but also to mechanisms for

building information-processors, including brains, nervous systems, and chemical information processors of many sorts.

(E.g. digestion uses informed disassembly of complex structures to find valuable parts to be transported and used or stored elsewhere.)

So far, in answer to Bell (quoted above), I have tried to describe the need for evolutionary selection mechanisms to be supported by enabling mechanisms.

Others have noticed the problem denied by Bell, e.g. Kirschner and Gerhart added some important biological details to the theory of evolved construction-kits, though not (as far as I can tell) the ideas (e.g. about abstraction and parametrisation) presented in this paper.

Work by Ganti and Kauffman is also relevant.

-- and probably others unknown to me!

BIOLOGICAL USES OF ABSTRACTION

As organisms grow in size, weight and strength, the forces and torques required at joints and at contact points with other objects change.

So the genome needs to use the same design with changing forces depending on tasks. Special cases include forces needed to move and manipulate the torso, limbs, gaze direction, chewed objects, etc. 'Hard-wiring' of useful evolved control functions with mathematical properties can be avoided by using designs that allow changeable parameters -- a strategy frequently used by human programmers.

Such parametrisation can both allow for changes in size and shape of the organism as it develops, and for many accidentally discovered biologically useful abstractions that can be parametrised in such designs -- e.g. allowing the same mechanism to be used for control of muscular forces at different stages of development, with changing weights, sizes, moments of inertia, etc.

Even more spectacular generalisation is achievable by re-use of evolved construction-kits

-- not only across developmental stages of individuals within a species,

-- but also across different species that share underlying physical parametrised design patterns,

-- with details that vary between species sharing the patterns

(as in vertebrates, or the more specialised variations among primates, or among birds, or fish species).

Such shared design patterns across species can result either from species having common ancestry or from convergent evolution ’driven’ by common features of the environment,

e.g. re-invention of visual processing mechanisms might be driven by aspects of spatial structures and processes common to all locations on the planet, despite the huge diversity of contents.

Such use of abstraction to achieve powerful re-usable design features across different application domains is familiar to engineers, including computer systems engineers.

’Design sharing’ explains why the tree of evolution has many branch points, instead of everything having to evolve from one common root node.

Symbiosis also allows combination of separately evolved features.

Similar ’structure-sharing’ often produces enormous reductions in search-spaces in AI systems.

It is also common in mathematics: most proofs build on a previously agreed framework of concepts, formalisms, axioms, rules, and previously proved theorems. They don’t all start from some fundamental shared axioms.

If re-usable abstractions can be encoded in suitable formalisms (with different application-specific parameters provided in different design contexts), they can enormously speed up evolution of diverse designs for functioning organisms.

This is partly analogous to the use of memo-functions in software design (i.e. functions that store computed values so that they don’t have to be re-computed whenever required, speeding up computations enormously, e.g. in the Fibonacci function).

Another type of re-use occurs in (unfortunately named) ’object-oriented’ programming paradigms that use hierarchies of powerful re-usable design abstractions, that can be instantiated differently in different combinations, to meet different sets of constraints in different environments, without requiring each such solution to be coded from scratch: ’parametric polymorphism’ with multiple inheritance.

This is an important aspect of many biological mechanisms. For example, there is enormous variation in what information perceptual mechanisms acquire and how the information is processed, encoded, stored, used, and in some cases communicated.

But abstract commonalities of function and mechanism (e.g. use of wings) can be combined with species specific constraints (parameters).

Parametric polymorphism makes the concept of consciousness difficult to analyse: there are many variants depending on what sort of thing is conscious, what it is conscious of, what information is acquired, what mechanisms are used, how the information contents are encoded, how they are accessed, how they are used, etc.

MATHEMATICAL CONSCIOUSNESS

Mathematical consciousness, still missing from AI, requires awareness of possibilities and impossibilities not restricted to particular objects, places or times -- as Kant pointed out.

Mechanisms and functions with mathematical aspects are also shared across groups of species, such as phototropism in plants, use of two eyes with lenses focused on a retina in many vertebrates, a subset of which evolved mechanisms using binocular disparity for 3-D perception.

That's one of many implicit mathematical discoveries in evolved designs for spatio-temporal perceptual, control and reasoning mechanisms, using the fact that many forms of animal perception and action occur in 3D space plus time, a fact that must have helped to drive evolution of mechanisms for representing and reasoning about 2-D and 3-D structures and processes, as in Euclidean geometry.

In a search for effective designs, enormous advantages come from (explicit or implicit) discovery and use of mathematical abstractions that are applicable across different designs or different instances of one design.

For example a common type of grammar (e.g. a phrase structure grammar) allows many different languages to be implemented including sentence generators and sentence analysers re-using the same program code with different grammatical rules.

Evolution seems to have discovered something like this.

Likewise, a common design framework for flying animals may allow tradeoffs between stability and maneouvreability to be used to adapt to different environmental opportunities and challenges.

These are mathematical discoveries implicitly used by evolution.

Evolution's ability to use these discoveries depends in part on the continual evolution of new DCKs providing materials, tools, and principles that can be used in solving many design and manufacture problems.

In recently evolved species, individuals e.g. humans and other intelligent animals, are able to replicate some of evolution's mathematical discoveries and make practical use of them in their own intentions, plans and design decisions, far more quickly than natural selection could.

Only (adult) humans seem to be aware of doing this.

Re-usable inherited abstractions allow different collections of members of one species, (e.g. humans living in deserts, in jungles, on mountain ranges, in arctic regions, etc.) to acquire expertise suited to their particular environments in a much shorter time than evolution would have required to produce the same variety of packaged competences 'bottom up'.

This flexibility also allows particular groups to adapt to major changes in a much shorter time than adaptation by natural selection would have required. This requires some later developments in individuals to be delayed until uses of earlier developments have provided enough information about environmental features to influence the ways in which later developments occur, as explained later.

This process is substantially enhanced by evolution of metacognitive information processing mechanisms that allow individuals to reflect on their own processes of perception, learning, reasoning, problem-solving, etc. and (to some extent) modify them to meet new conditions.

Later, more sophisticated products of evolution develop metameta-cognitive information processing sub-architectures that enable them to notice their own adaptive processes, and to reflect on and discuss what was going on, and in some cases collaboratively improve the processes,

-- e.g. through explicit teaching

-- at first in a limited social/cultural context, after which the activity was able to spread

-- using previously evolved learning mechanisms.

As far as I know only humans have achieved that, though some other species apparently have simpler variants.

These conjectures need far more research!

Human AI designs for intelligent machines created so far seem to have far fewer layers of abstraction, and are far more primitive, than the re-usable designs produced by evolution. Studying the differences is a major sub-task facing the M-M project (and AI).

This requires a deep understanding of what needs to be explained.

DESIGNING DESIGNS

Just as the designer of a programming language cannot know about, and does not need to know about, all the applications for which the programming language will be used, so also can the more abstract products of evolution be instantiated (e.g. by setting parameters) for use in contexts in which they did not evolve.

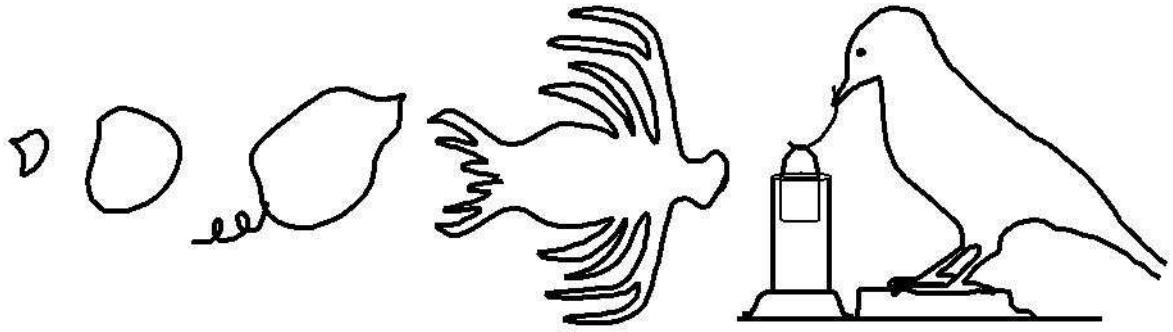

Many discontinuities in physical forms, behavioural capabilities, environments, types of information acquired, types of use of information and mechanisms for information-processing are still waiting to be discovered.

EVOLUTION OF HUMAN LANGUAGE CAPABILITIES

One of the most spectacular cases is reuse of a common collection of language-creation competences in a huge variety of geographical and social contexts, allowing any individual human to acquire any of several thousand enormously varied human languages, including both spoken and signed languages.

A striking example was the cooperative creation by deaf children in Nicaragua of a new sign language because their teachers had not learned sign languages early enough to develop full adult competences. This suggests that what is normally regarded as language learning is really cooperative language creation, demonstrated in this video:

https://www.youtube.com/watch?v=pjtioIFuNf8

Re-use can take different forms, including

-- re-use of a general design across different species by instantiating a common pattern,

-- re-use based on powerful mechanisms for acquiring and using information about the available resources, opportunities and challenges during the development of each individual.

The first process happens across evolutionary lineages.

The second happens within individual organisms in their lifetime

Social/cultural evolution requires intermediate timescales.

Evolution seems to have produced multi-level design patterns, whose details are filled in incrementally, during creation of instances of the patterns in individual members of a species.

If all the members live in similar environments that will tend to produce uniform end results.

However, if the genome is sufficiently abstract, then environments and genomic structures may interact in more complex ways, allowing small variations during development of individuals to cascade into significant differences in the adult organism, as if natural selection had been sped up enormously.

A special case is evolution of an immune system with the ability to develop different immune responses depending on the antigens encountered. Another dramatic special case is the recent dramatic cascade of social, economic, and educational changes supported jointly by the human genome and the internet!

CHANGES IN DEVELOPMENTAL TRAJECTORIES

As living things become more complex, increasingly varied types of information are required for increasingly varied uses.

The processes of reproduction normally produce new individuals that have seriously under-developed physical structures and behavioural competences.

Self-development requires physical materials, but it also requires information about what to do with the materials, including disassembling and reassembling chemical structures at a sub-microscopic level and using the products to assemble larger body parts, while constantly providing new materials, removing waste products and consuming energy.

Some energy is stored and some is used in assembly and other processes.

The earliest (simplest?) organisms can acquire and use information about (i.e. sense) only internal states and processes and the immediate external environment, e.g. pressure, temperature, and presence of chemicals in the surrounding soup, with all uses of information taking the form of immediate local reactions, e.g. allowing a molecule through a membrane.

Changes in types of information, types of use of information and types of biological mechanism for processing information have repeatedly altered the processes of evolutionary morphogenesis that produce such changes: a positive feedback process.

An example is the influence of mate selection on evolution in intelligent organisms: mate selection is itself dependent on previous evolution of cognitive mechanisms. Hence the prefix ’Meta-’ in ’Meta-Morphogenesis’.

This is a process with multiple feedback loops between new designs and new requirements (niches), as suggested in

ONLINE VS OFFLINE INTELLIGENCE

As the previous figure suggests, evolution constantly produces new organisms that may or may not be larger than predecessors, but are more complex both in the types of physical action they can produce and also the types of information and types of information processing required for selection and control of such actions.

Some of that information is used immediately and discarded (online perceptual intelligence) while other kinds are stored, possibly in transformed formats, and used later, possibly on many occasions (offline perceptual intelligence) -- a distinction often mislabelled as ’where’ vs ’what’ perception.

This generalises Gibson’s theory that perception mainly provides information about ’affordances’ rather than information about visible surfaces of perceived objects.

These ideas, like Karmiloff-Smith’s Beyond Modularity suggest that one of the effects of biological evolution was fairly recent production of more or less abstract construction kits that come into play at different stages in development, producing new more rapid changes in variety and complexity of information processing across generations as explained below (See fig 2)

It’s not clear how much longer this can continue: perhaps limitations of human brains constrain this process. But humans working with intelligent machines may be able to stretch the limits.

At some much later date, probably in another century, we may be able to make machines that do it all themselves -- unless it turns out that the fundamental information processing mechanisms in brains cannot be modelled in computer technology developed by humans.

Species can differ in the variety of types of sensory information they can acquire, in the variety of uses to which they put that information, in the variety of types of physical actions they can produce, in the extent to which they can combine perceptual and action processes to achieve novel purposes or solve novel problems, and the

extent to which they can educate, reason about, collaborate with, compete against conspecifics, and prey or competitor species.

As competences become more varied and complex, the more disembodied must the information processing be, i.e. disconnected from current sensory and motor signals (while preserving low level reflexes and sensory-motor control loops for special cases).

This may have been a precursor to mathematical abilities to think about transfinite set theory and high dimensional vector spaces or complex modern scientific theories.

E.g. Darwin's own thinking about ancient evolutionary processes. was detached from his particular sensory-motor processes at the time! This applies also to affective states, e.g. compare being startled and being obsessed with ambition.

The fashionable emphasis on "embodied cognition" may be appropriate to the study of organisms such as plants and microbes, and perhaps insects, but evolved intelligence increasingly used disembodied cognition, most strikingly in the production of ancient mathematical minds. This led to new complexities in processes of epigenesis (gene-influenced development).

Waddington's view of epigenesis

A ball rolling (passively) down a fixed landscape

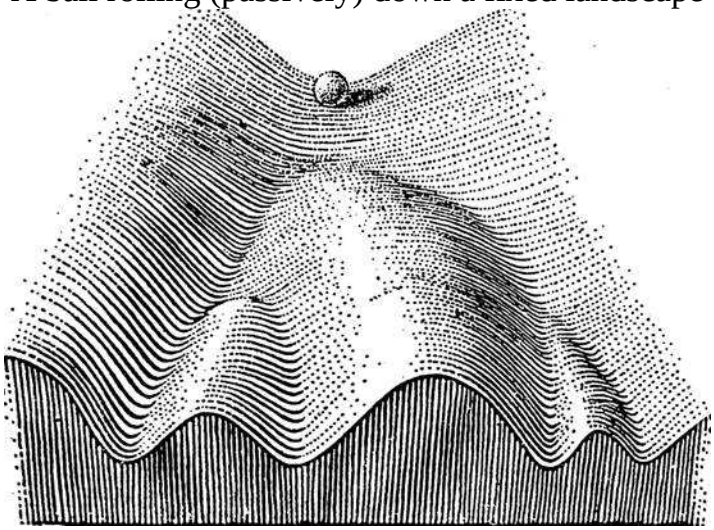

Figure WAD:

A more recent picture of epigenesis (beyond Waddington)

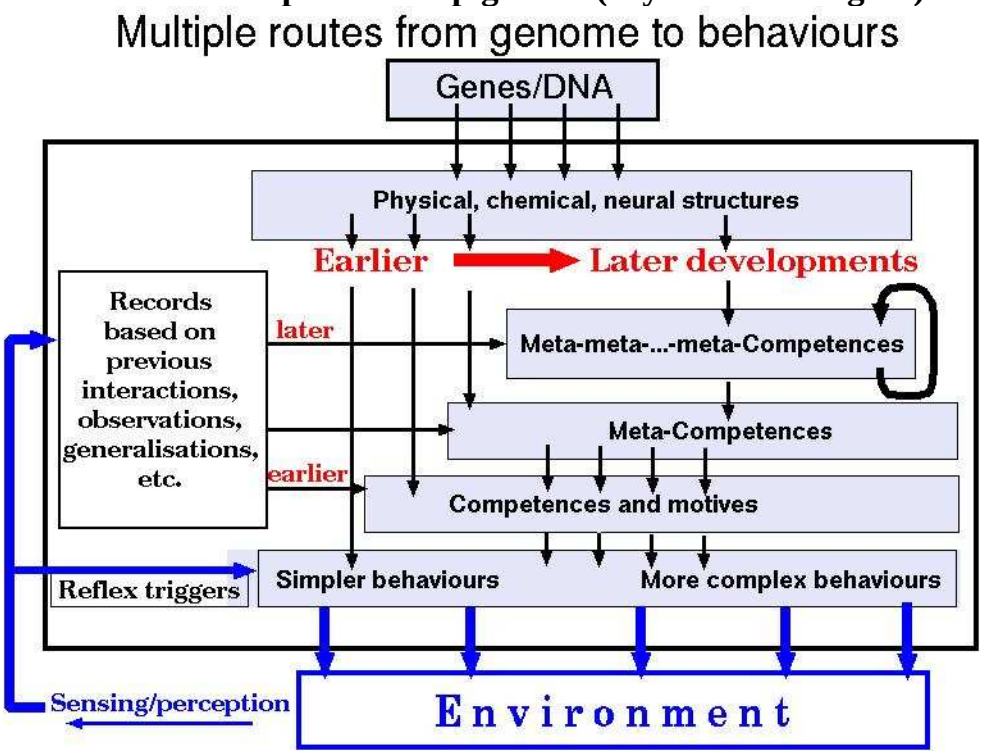

Figure EPI:
Cascaded, staggered, developmental trajectories, with later processes influenced by results of earlier processes in increasingly complex ways. Proposed by Chappell and Sloman 2007[3]

Early genome-driven learning from the environment occurs in loops on the left. Downward arrows further right represent later gene-triggered processes during individual development modulated by results of earlier learning via feedback on left.

(Chris Miall suggested the structure of the original diagram.)

VARIATIONS IN EPIGENETIC TRAJECTORIES

The description given so far is very abstract and allows significantly different instantiations in different species, addressing different sorts of functionality and different types of design, e.g. of physical forms, behaviours, control mechanisms, reproductive mechanisms, etc.

At one extreme the reproductive process produces individuals whose genome exercises a fixed pattern of control during development, leading to ’adults’ with only minor variations.

At another extreme, instead of the process of development from one stage to another being fixed in the genome, it could be created during development through the use of

more than one level of design in the genome.

E.g. if there are two levels then results of environmental interaction at the first level could transform what happens at the second level. If there are multiple levels then what happens at each new level may be influenced by results of earlier developments.

In a species with such multi-stage development, at intermediate stages not only are there different developmental trajectories due to different environmental influences, there are also selections among the intermediate level patterns to be instantiated, so that in one environment development may include much learning concerned with protection from freezing, whereas in other environments individual species may vary more in the ways they seek water during dry seasons.

Then differences in adults come partly from the influence of the environment in selecting patterns to instantiate. E.g. one group may learn and pass on information about where the main water holes are, and in another group individuals may learn and pass on information about which plants are good sources of water.

If these conjectures are correct, patterns of development will automatically be varied because of patterns and meta-patterns picked up by earlier generations and instantiated in cascades during individual development.

So different cultures produced jointly by a genome and previous environments can produce very different expressions of the same genome, even though individuals share similar physical forms.

The main differences are in the kinds of information acquired and used, and the information processing mechanisms developed. Not all cultures use advanced mathematics in designing buildings, but all build on previously evolved understanding of space, time and motion.

Evolution seems to have found how to provide rich developmental variation by allowing information gathered by young individuals not merely to select and use pre-stored design patterns, but to create new patterns by assembling fragments of information during earlier development, then using more abstract processes to construct new abstract patterns, partly shaped by the current environment, but with the power to be used in new environments.

Developments in culture (including language, science, engineering, mathematics, music, literature, etc.) all show such combinations of data collection and enormous creativity, including creative ontology extension (e.g. the Nicaraguan children mentioned above.

Unless I have misunderstood her, this is the type of process Karmiloff-Smith called ’Representational Re-description’ (RR).

Genome-encoded previously acquired abstractions ’wait’ to be instantiated at different stages of development, using cascading alternations between data-collection and abstraction formation (RR) by instantiating higher level generative abstractions (e.g. meta-grammars), not by forming statistical generalisations.

This could account for both the great diversity of human languages and cultures, and the power of each one, all supported by a common genome operating in very different environments.

Jackie Chappell noticed the implication that instead of the genome specifying a fixed ’epigenetic landscape’ (proposed by Waddington) it provides a schematic landscape and mechanisms that allow each individual (or in same cases groups of individuals) to modify the landscape while moving down it (e.g. adding new hills, valleys, channels and barriers).

Though most visible in language development, the process is not unique to language development, but occurs throughout childhood (and beyond) in connection with many aspects of development of information processing abilities, construction of new ontologies, theory formation, etc.

This differs from forms of learning or development that use uniform statistics-based methods for repeatedly finding patterns at different levels of abstraction.

Instead, Figure 2 indicates that the genome encodes increasingly abstract and powerful creative mechanisms developed at different stages of evolution, that are ’awakened’ (a notion used by Kant) in individuals only when appropriate, so that they can build on what has already been learned or created in a manner that is tailored to the current environment.

For example, in young (non-deaf) humans, processes giving sound sequences a syntactic interpretation develop after the child has learnt to produce and to distinguish some of the actual speech sounds used in that location.

In social species, the later stages of Figure 2 include mechanisms for discovering non-linguistic ontologies and facts that older members of the community have acquired, and incorporating relevant subsets in combination with new individually acquired information.

Instead of merely absorbing the details of what older members have learnt, the young can absorb forms of creative learning, reasoning and representation that older members have found useful and apply them in new environments to produce new results.

In humans, this has produced spectacular effects, especially in the last few decades.

The evolved mechanisms for representing and reasoning about possibilities, impossibilities and necessities were essential for both perception and use of affordances and for making mathematical discoveries, something statistical learning cannot achieve.

SPACE-TIME

An invariant for all species in this universe is space-time embedding, and changing spatial relationships between body parts and things in the environment.

The relationships vary between water-dwellers, cave-dwellers, tree-dwellers, flying animals, and modern city-dwellers.

Representational requirements depend on body parts and their controllable relationships to one another and other objects.

So aeons of evolution will produce neither a tabula rasa nor geographically specific spatial information, but a collection of generic mechanisms for finding out what sorts of spatial structures have been bequeathed by ancestors as well as physics and geography, and learning to make use of whatever is available (McCarthy[17]): that's why embodiment is relevant to evolved cognition.

Kant's ideas about geometric knowledge are relevant though he assumed that the innate apparatus was geared only to structures in Euclidean space, whereas our space is only approximately Euclidean.

Somehow the mechanisms conjectured in Figure 2 eventually (after many generations) made it possible for humans to make the amazing discoveries recorded in Euclid's Elements, still used world-wide by scientists and engineers.

If we remove the parallel axiom we are left with a very rich collection of facts about space and time, especially topological facts about varieties of structural change, e.g. formation of networks of relationships, deformations of surfaces, and possible trajectories constrained by fixed obstacles.

It is well known (though non-trivial to prove!) that trisection of an arbitrary angle is impossible in Euclidean geometry, whereas bisection is trivial.

However, some ancient mathematicians (e.g. Archimedes) knew that there is a fairly simple addition to Euclidean geometry that makes trisecting an arbitrary angle easy, namely the 'neusis' construction that allows a movable straight edge to have two marks fixed on it that can be used to specify constraints on motion of the edge. http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html

They proved this without modern logic, algebra, set theory, proof theory etc. However, there is no current AI reasoner capable of discovering such a construct, or considering whether it is an acceptable extension to Euclid's straight-edge and compasses constructs.

If we can identify a type of construction-kit that produces young robot minds able to develop or evaluate those ideas in varied spatial environments, we may find important clues about what is missing in current AI.

Long before logical and algebraic notations were used in mathematical proofs, evolution had produced abilities to represent and reason about what Gibson called 'affordances', including possible and impossible alterations to spatial configurations

Example:

> The (topological) impossibility of solid linked rings becoming unlinked, or vice versa.
> See also this rubber-band example:
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/rubber-bands.html

I suspect brains of many intelligent animals make use of topological reasoning mechanisms that have so far not been discovered by brain scientists or AI researchers.

Addition of meta-cognitive mechanisms able to inspect and experiment with reasoning processes may have led both to enhanced spatial intelligence and meta-cognition, and also to meta-metacognitive reasoning about other intelligent individuals.

OTHER SPECIES

I conjecture that further investigation will reveal varieties of information processing (computation) that have so far escaped the attention of researchers, but which play important roles in many intelligent species, including not only humans and apes but also elephants, corvids, squirrels, cetaceans and others.

In particular, some intelligent non-human animals and pre-verbal human toddlers seem to be able to use mathematical structures and relationships (e.g. partial orderings and topological relationships) unwittingly. Mathematical meta-meta...-cognition seems to be restricted to humans, but develops in stages, as Piaget found, partially confirming Kant's ideas about mathematical knowledge in.

However, I suspect that (as Kant seems to have realised) the genetically provided mathematical powers of intelligent animals make more use of topological and geometric reasoning, using analogical, non-Fregean, representations, as suggested in

than the logical, algebraic, and statistical capabilities that have so far dominated AI and robotics.

For example, even the concepts of cardinal and ordinal number are crucially related to concepts of one-one correspondence between components of structures, most naturally understood as a topological relationship rather than a logically definable relationship. http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap8.html

(NB 'analogical' does not imply 'isomorphic' as often suggested. A typical 2D picture (an analogical representation) of a 3D scene cannot be isomorphic with the scene depicted. A projection is not an isomorphism if it removes some of the relationships. There is a deeper distinction between Fregean and Analogical forms of representation Sloman (1971), concerned with the relationships between representation and what is represented.

DISEMBODIMENT OF COGNITION EVOLVES

All this shows why increasing complexity of physical structures and capabilities, providing richer collections of alternatives and more complex internal and external action-selection criteria, requires increasing disembodiment of information processing.

The fact that evolution is not stuck with the Fundamental Construction Kit (FCK) provided by physics and chemistry, but also produces and uses new 'derived' construction-kits (DCKs), enhances both the mathematical and the ontological creativity of evolution, which is indirectly responsible for all the other known types of creativity.

This counters both the view that mathematics is a product of human minds, and a view of metaphysics as being concerned with something unchangeable.

The notion of 'Descriptive Metaphysics' presented by Strawson (1959) needs to be revised.

DO WE NEED NON-TURING FORMS OF COMPUTATION?

I also conjecture that filling in some of the missing details in this theory (a huge challenge) will help us understand both the evolutionary changes that introduced unique features of human minds and why it is not obvious that Turing-equivalent digital computers, or even asynchronous networks of such computers running sophisticated interacting virtual machines, will suffice to replicate the human mathematical capabilities that preceded modern logic, algebra, set-theory, and theory of computation.

It will all depend on the precise forms of virtual information processing machinery that evolution has managed to produce, about which I suspect current methods of neuroscientific investigation cannot yield deep information.

Current AI cannot produce reasoners like Euclid, Zeno, Archimedes, or even reasoners like pre-verbal toddlers, weaver birds and squirrels.

This indicates serious gaps, despite many impressive achievements. I see no reason to believe that uniform, statistics based learning mechanisms will have the power to bridge those gaps.

WHAT ABOUT LOGIC?

Whether the addition of logic-based reasoners will suffice (as suggested by McCarthy and Hayes)(1969) is not clear.

The discoveries made by ancient mathematicians preceded the discoveries of modern algebra and logic, and the arithmetisation of geometry by Descartes.

Evolved mechanisms that use previously acquired abstract forms of meta-learning with genetically orchestrated instantiation triggered by developmental changes (as in the above diagram), may do much better.

These mechanisms depend on rich internal languages that evolved for use in perception, reasoning, learning, intention formation, plan formation and control of actions before communicative languages.

This generalises claims made by Chomsky in, and his later works, focused only on development of human spoken languages, ignoring how much language and non-linguistic cognition develop with mutual support.

THE IMPORTANCE OF VIRTUAL MACHINERY

Building a new computer for every task was made unnecessary by allowing computers to have changeable programs.

Initially each program, specifying instructions to be run, had to be loaded (via modified wiring, switch settings, punched cards, or punched tape), but later developments provided more and more flexibility and generality, with higher level programming languages providing reusable domain specific languages and tools, some translated to machine code, others run on a task specific virtual computer provided by an interpreter.

Later developments provided time-sharing operating systems supporting multiple interacting programs running effectively in parallel performing different, interacting,

tasks on a single processor.

As networks developed, these collaborating virtual machines became more numerous, more varied, more geographically distributed, and more sophisticated in their functionality, often extended with sensors of different kinds and attached devices for manipulation, carrying, moving, and communicating.

These developments suggest the possibility that each biological mind is also implemented as a collection of concurrently active nonphysical, but physically implemented, virtual machines interacting with one another and with the physical environment through sensor and motor interfaces.

Such ’virtual machine functionalism’ could accommodate a large variety of coexisting, interacting, cognitive, motivational and emotional states, including essentially private qualia as explained by Sloman and Chrisley (2003).

Long before human engineers produced such designs, biological evolution had already encountered the need and produced virtual machinery of even greater complexity and sophistication, serving information processing requirements for organisms, whose virtual machinery included interacting sensory qualia, motivations, intentions, plans, emotions, attitudes, preferences, learning processes, and various aspects of self-consciousness.

THE FUTURE OF AI

We still don’t know how to make machines able to replicate the mathematical insights of ancient mathematicians like Euclid e.g. with ’triangle qualia’ that include awareness of mathematical possibilities and constraints, or minds that can discover the possibility of extending Euclidean geometry with the neusis construction. For discussion of roles of ’triangle qualia’ in discoveries made by ancient mathematicians see
triangle-theorem.html http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ triangle-theorem.html
The use of the "neusis" construction to trisect an arbitrary angle is explained in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/neusis.html

NOTE
It is not clear whether we simply have not been clever enough at understanding the problems and developing the programs, or whether we need to extend the class of virtual machines that can be run on computers, or whether the problem is that animal brains use kinds of virtual machinery that cannot be implemented using the construction kits known to modern computer science and software engineering. As Turing hinted in his 1950 paper: aspects of chemical computation may be essential.

Biological organisms also cannot build such minds directly from atoms and molecules. They need many intermediate DCKs, some of them concrete and some abstract, insofar as some construction kits, like some animal minds, use virtual machines.

Evolutionary processes must have produced construction kits for abstract information processing machinery supporting increasingly complex multi-functional virtual machines, long before human engineers discovered the need for such things and began to implement them in the 20th Century.

Studying such processes is very difficult because virtual machines don't leave fossils (though some of their products do). Moreover details of recently evolved virtual machinery may be at least as hard to inspect as running software systems without built-in run-time debugging 'hooks'. This could, in principle, defeat all known brain scanners.

'Information' here is not used in Shannon's sense (concerned with mechanisms and vehicles for storage, encoding, transmission, decoding, etc.), but in the much older sense familiar to Jane Austen and used in her novels e.g. Pride and Prejudice, in which how information content is used is important, not how information bearers are encoded, stored, transmitted, received, etc. The primary use of information is for control.

Communication, storage, reorganisation, compression, encryption, translation, and many other ways of dealing with information are all secondary to the use for control. Long before humans used structured languages for communication, intelligent animals must have used rich languages with structural variability and compositional semantics internally, e.g. in perception, reasoning, intention formation, wondering whether, planning and execution of actions, and learning.

We can search for previously unnoticed evolutionary transitions going beyond the examples here (e.g. Figure 1), e.g. transitions between organisms that merely react to immediate chemical environments in a primaeval soup, and organisms that use temporal information about changing concentrations in deciding whether to move or not.

Another class of examples seems to be the new mechanisms required after the transition from a liquid based life form to life on a surface with more stable structures (e.g. different static resources and obstacles in different places), or a later transition to hunting down and eating mobile land-based prey, or transitions to reproductive mechanisms requiring young to be cared for, etc.? Perhaps we'll then understand how to significantly extend AI.

Compare Schrödinger's discussion in [19] of the relevance of quantum mechanisms and chemistry to the storage, copying, and processing of genetic information.26 I am suggesting that questions about evolved intermediate forms of information processing are linked to philosophical questions about the nature of mind, the nature of mathematical discovery, and deep gaps in current AI.27

NOTES:

19 Boden [2] distinguishes H-Creativity, which involves being historically original, and P-Creativity, which requires only personal originality. The distinction is echoed in the phenomenon of convergent evolution, illustrated in https://en.wikipedia.org/wiki/List%20of%20examples%20of%20convergent%20evolution The first species with some design solution exhibits H-creativity of evolution. Species in which that solution evolves independently later exhibit a form of P-creativity.

20 Why did Turing write in his his 1950 paper that chemistry may turn out to be as important as electricity in brains?

REFERENCES
To be re-formatted, with links.

[1] Graham Bell, *Selection The Mechanism of Evolution,* OUP, 2008. Second Edition.

[2] M. A. Boden, *The Creative Mind: Myths and Mechanisms,* Weidenfeld & Nicolson, London, 1990. (Second edition, Routledge, 2004).

[3] Jackie Chappell and Aaron Sloman, "Natural and artificial metaconfigured altricial information-processing systems", *International Journal of Unconventional Computing,* 3(3), 221-239, (2007).

[4] N. Chomsky, *Aspects of the theory of syntax,* MIT Press, Cambridge, MA, 1965.

[5] Juliet C. Coates, Laura A. Moody, and Younousse Saidi, "Plants and the Earth system - past events and future challenges', *New Phytologist,* 189, 370-373, (2011).

[6] *Alan Turing - His Work and Impact,* eds., S. B. Cooper and J. van Leeuwen, Elsevier, Amsterdam, 2013. (contents list).

[7] T. Froese, N. Virgo, and T. Ikegami, Motility at the origin of life: Its characterization and a model', *Artificial Life,* 20(1), 55-76, (2014).

[8] Tibor Ganti, 2003 *The Principles of Life,* OUP, New York, Eds. Eors Szathmary & James Griesemer, Translation of the 1971 Hungarian edition.

[9] J. J. Gibson, *The Ecological Approach to Visual Perception,* Houghton Mifflin, Boston, MA, 1979.

[10] M. M. Hanczyc and T. Ikegami, 'Chemical basis for minimal cognition', *Artificial Life,* 16, 233-243, (2010).

[11] John Heslop-Harrison, *New concepts in flowering-plant taxonomy*, Heinemann, London, 1953.

[12] Immanuel Kant, *Critique of Pure Reason*, Macmillan, London, 1781. Translated (1929) by Norman Kemp Smith.
Various online versions are also available now.

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

[14] S. Kauffman, *At home in the universe: The search for laws of complexity*, Penguin Books, London, 1995.

[15] M.W. Kirschner and J.C. Gerhart, *The Plausibility of Life: Resolving Darwin's Dilemma*, Yale University Press, Princeton, 2005.

[16] D. Kirsh, "Today the earwig, tomorrow man?', *Artificial Intelligence*, 47(1), 161-184, (1991).

I. Lakatos, 1976, *Proofs and Refutations*, Cambridge University Press, Cambridge, UK,

[17a] John McCarthy and Patrick J. Hayes, 1969, "Some philosophical problems from the standpoint of AI", *Machine Intelligence 4*, Eds. B. Meltzer and D. Michie, pp. 463--502, Edinburgh University Press,
http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html

[17] J. McCarthy, "The well-designed child', *Artificial Intelligence*, 172(18), 2003-2014, (2008).

[18] W. T. Powers, *Behavior, the Control of Perception*, Aldine de Gruyter, New York, 1973.

[19] Erwin Schrödinger, *What is life?*, CUP, Cambridge, 1944.
Commented extracts available here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html

[20] A. Sloman, 1962, *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth* (DPhil Thesis), PhD. dissertation, Oxford University, (now online)
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-1962

[21] A. Sloman, 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. Reprinted in *Artificial Intelligence*, vol 2, 3-4, pp 209-225, 1971.
http://www.cs.bham.ac.uk/research/cogaff/62-80.html#1971-02
An expanded version was published as chapter 7 of Sloman 1978, available here.

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

[23] A. Sloman, (2000) "Interacting trajectories in design space and niche space: A philosopher speculates about evolution', in *Parallel Problem Solving from Nature (PPSN VI)*, eds. M.Schoenauer, et al. Lecture Notes in Computer Science, No 1917, pp. 3-16, Berlin, (2000). Springer-Verlag.

[24] A. Sloman and R.L. Chrisley, (2003) "Virtual machines and consciousness', *Journal of Consciousness Studies*, 10(4-5), 113-172.

[25] Aaron Sloman,2013a "Virtual Machine Functionalism (The only form of functionalism worth taking seriously in Philosophy of Mind and theories of Consciousness)', Research note, School of Computer Science, The University of Birmingham.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-func.html

[26] Aaron Sloman,2013b "Virtual machinery and evolution of mind (part 3) Meta-morphogenesis: Evolution of information-processing machinery', in *Alan Turing - His Work and Impact*, eds., S. B. Cooper and J. van Leeuwen, 849-856, Elsevier, Amsterdam.
http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d

[27] Aaron Sloman (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

[27a] Aaron Sloman 2017, "Construction kits for evolving life (Including evolving minds and mathematical abilities.)" Technical report (work in progress)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

An earlier version, frozen during 2016, was published in a Springer Collection in 2017:
https://link.springer.com/chapter/10.1007%2F978-3-319-43669-2_14
in *The Incomputable Journeys Beyond the Turing Barrier*
Eds: S. Barry Cooper and Mariya I. Soskova
https://link.springer.com/book/10.1007/978-3-319-43669-2

[28] Aaron Sloman and David Vernon. A First Draft Analysis of some MetaRequirements for Cognitive Systems in Robots, 2007. Contribution to euCognition wiki.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-requirements.html

[29] P. F. Strawson, *Individuals: An essay in descriptive metaphysics*, Methuen, London, 1959.

[29a] Max Tegmark, 2014, *Our mathematical universe, my quest for the ultimate nature of reality*, Knopf (USA) Allen Lane (UK), (ISBN 978-0307599803/978-1846144769)

[30] A. M. Turing, "Computing machinery and intelligence', *Mind*, 59, 433-460, (1950). (reprinted in E.A. Feigenbaum and J. Feldman (eds) Computers and Thought McGraw-Hill, New York, 1963, 11-35).

[31] A. M. Turing, "The Chemical Basis Of Morphogenesis", *Phil. Trans. Royal Soc. London B* 237, 237, 37-72, (1952).

Note: A presentation of Turing's main ideas for non-mathematicians can be found in
Philip Ball, 2015, "Forging patterns and making waves from biology to geology: a commentary on Turing (1952) 'The chemical basis of morphogenesis'",
http://dx.doi.org/10.1098/rstb.2014.0218

[32] C. H. Waddington, *The Strategy of the Genes. A Discussion of Some Aspects of Theoretical Biology*, George Allen & Unwin, 1957.

[33] R. A. Watson and E. Szathmary, "How can evolution learn?', *Trends in Ecology and Evolution*, 31(2), 147-157, (2016).

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