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

(Or several levels of information-processing machinery)

DRAFT DRAFT DRAFT DRAFT

EVOLVED LEVELS IN OUR UNIVERSE
Universe-level Epigenetics

(DRAFT: Liable to change)

ICCM Warwick University July 2017:
http://mathpsych.org/conferences/2017/file/MP_ICCM2017_Program.pdf
PROCEEDINGS:
http://iccm-conference.org/2017/
(Embedded in Programme)

Notes related to presentation on Monday 24th July

Gaps Between Human and (Current) Artificial Mathematics
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman_iccm_17.pdf

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

This is part of the Meta-Morphogenesis (M-M) project
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html

Including the theory of evolved construction kits
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

---

This paper is http://www.cs.bham.ac.uk/research/projects/cogaff/misc/levels.html
Abbreviated link to this file http://goo.gl/TRG6ZF
Also PDF:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/levels.html

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

---

Abstract published for ICCM
The Turing-inspired Meta-morphogenesis project begun in 2011 was partly motivated by deep gaps in our understanding of mathematical cognition and other aspects of human and non-human intelligence and our inability to model them. The project attempts to identify previously unnoticed evolutionary transitions in biological information processing related to gaps in our current understanding of cognition. Analysis of such transitions may also shed light on gaps in current AI. This is very different from attempts to study human mathematical cognition directly, e.g. via observation, experiment, neural imaging, etc. Fashionable ideas about "embodied cognition", "enactivism", and "situated cognition", focus on shallow products of evolution, ignoring pressures to evolve increasingly disembodied forms of cognition to meet increasingly complex and varied challenges produced by articulated physical forms, multiple sensory capabilities, geographical and temporal spread of important information and other resources, and "other-related meta-cognition" concerning mental states, processes and capabilities of other individuals. Computers are normally thought of as good at mathematics: they perform logical, arithmetical and statistical calculations and manipulate formulas, at enormous speeds, but still lack abilities in humans and other animals to perceive and understand geometrical and topological possibilities and constraints that (a) are required for perception and use of affordances, and (b) play roles in mathematical, and proto-mathematical, discoveries made by ancient mathematicians, human toddlers and other intelligent animals. Neurally inspired, statistics-based (e.g."deep learning") models cannot explain recognition and understanding of mathematical necessity or impossibility. A partial (neo-Kantian) analysis of types of evolved biological information processing capability still missing from our models may inspire new kinds of research helping to fill the gaps. Had Turing lived long enough to develop his ideas on morphogenesis, he might have done this.

Keywords:
Thales, Archimedes; Euclid; Kant; geometry; topology; cognition; evolution; development; abstraction; non-empirical; non-contingent; toddler-theorems; disembodiment;

Submitted paper, with references and web links:
GAPS BETWEEN HUMAN and ARTIFICIAL MATHEMATICS
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman_iccm_17.pdf
Search for Aaron Sloman, ICCM 2017

ONE OF THE HARDEST PROBLEMS IN SCIENCE:

Understanding natural intelligence

Including squirrel intelligence, crow intelligence, elephant intelligence, human toddler intelligence, octopus intelligence, Portia spider intelligence, ...

I am specially interested in intelligence of ancient mathematicians who lived over two thousand years ago, e.g. Archimedes, Euclid, Pythagoras, Thales ... and their predecessors.

Discoveries/inventions of such ancient mathematicians are still in use by scientists and engineers all over the planet every day.

Part of our problem is to describe how these discoveries differ from others.

Mathematical truths can be discovered empirically, but there is also another way: proving them to be truths. Empirical discovery leaves open the possibility of an exception turning up. That's not possible necessary truths.

If you have proved the triangle sum theorem (interior angles of a triangle add up to a straight line in a Euclidean plane, though not on a sphere or torus, etc.), then you don't need to keep experimenting with triangles of different colours, oriented in different directions, at different altitudes, or temperatures, etc.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html

Current AI models/explanations don't explain what made those discoveries possible long before the discovery (by Descartes and others) that Euclidean geometry can be modelled using arithmetic, algebra and logic. Geometrical reasoning by ancient mathematicians used diagrams (real or imagined) not equations and logic.

Neither (as far as I know) do psychology, neuroscience, philosophy, provide adequate explanations of ancient mathematical reasoning in a form that can be replicated in computers or robots.

Moreover, most researchers don't even have adequate specifications of what needs to be explained -- they have not read Kant! (More on Kant below.)

Current AI systems can attach labels to spatial structures, and can learn correlations between them, including consequences of drawing additional lines or circles. But they cannot (yet) grasp that some construction is impossible, or that one changed relationship in a spatial configuration necessarily produces some other change.

For example, if a vertex of a planar triangle is moved further from the opposite side of the triangle, the angle at that vertex must get smaller. Why?

Similarly, a child who understands the role of one-to-one correspondences in the game of musical chairs can reason about the consequences of a change in the way the game is played, e.g. arranging the chairs in a circle instead of a line, or removing two chairs at a time instead of only one. (How does that change the possible outcomes?) What about removing three at a time? (These questions require abilities to reason about properties of one-to-one correspondences.)

Neither AI, nor psychology, nor neuroscience sheds light on what is involved in the process of noticing that something is a **necessary** truth or that something is **necessarily** false.

> Mathematical discoveries are not empirical discoveries about what occurs or is probable in various conditions.
> As Immanuel Kant pointed out in 1781, they are about necessary truths, impossibilities, and necessary consequences, and these cannot be demonstrated by collecting statistical information, and using probabilistic reasoning (even forms of "deep" learning).
>
> A difference between mathematical necessary truths and contingent truths is that no counter example to the former can be found or constructed though that is possible for contingent truths, though it may be very difficult to do.

A logical or algebraic reasoning system can establish such necessary truths by exhaustively testing permitted changes of symbols.

Humans can also discover some truths by exhaustive analysis of possible continuous deformations of geometrical structures, in ways that current automated reasoning systems cannot, e.g. discovering that a closed curve C1 on the surface of a torus cannot be continuously deformed into another curve C2 in a different location on the same surface, as illustrated in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html

There have been computer-based "Geometry theorem provers" since the earliest days of AI, e.g. Gelernter(1964) (apparently inspired by some pencil and paper experiments by Minsky (see Bernstein Interview (1981)) but those geometrical theorem proving programs start from Euclid's axioms expressed in a logical notation, from which logically valid inferences are used to derive the desired conclusion.

> Sometimes conjectured (proof-subgoals) are eliminated by constructing and examining "models" that turn out to be counter examples. Humans can do this by constructing or visualising a new geometric configuration. AI theorem provers that cannot visualise instead use diagrams specified using numerical coordinates, which sometimes speed up proofs by refuting conjectured intermediate goals, avoiding time wasted trying to prove false sub-goals.
>
> Following Descartes, this automated proof-pruning mechanism uses numbers as coordinates of points, and also as properties of lines (e.g. slope and intercept). That arithmetical "translation" has only been known since the time of Descartes, hence the label "cartesian coordinates". It could not have been used by the ancient mathematicians.
>
> Despite the extent to which arithmetic can model spatial structures composed of straight lines, circles and ellipses, computers can put the correspondence to use. But using such numerical models is a very different process from using the visualised or drawn geometric or topological models available to ancient mathematicians.

Unlike such automated theorem provers, Euclid and his predecessors did NOT start their discoveries by assuming the axioms in Euclid's *Elements*. Rather those axioms, e.g. for congruence of triangles, were originally *discovered* to be true of triangles, though exactly how is not known.

The axioms in Euclid's *Elements* were major discoveries arising out of features of human (and some animal) forms of consciousness of spatial structures and processes. (I.e. the axioms were not just arbitrary implicit definitions of the subject matter.)

So no reasoning model that starts from those, or any other logically expressed collection of axioms, can be a model of the original types of mathematical discovery.

Moreover, since (as Kant pointed out) they were not discoveries of contingent, *empirical* generalisations, nor analytic truths derivable from definitions by logic, but discoveries of necessary synthetic truths, all the current probability+statistics based learning mechanisms in AI (e.g. use of "deep learning" from "Big Data" is totally irrelevant to modelling ancient mathematics.

Moreover even the concepts used are not all abstracted from experience of instances, as many empirical concepts are. For example, it is impossible to experience an infinitely thin line, or an infinitely small point. Both are, at most, limiting cases of what can be experienced. Other examples include perfect straightness, perfect circularity, perfect planarity, perfect parallelism, etc.

As far as I know there is no concept formation method in AI, and no brain mechanism known to neuroscientists, and no form of representation in AI or neuroscience, that allows the processes by which these concepts were represented in brain mechanisms of ancient mathematicians to be modelled or explained by current AI programs or neural simulation models. How do brains represent infinitely long lines, infinite planes, etc.

As far as I know, there are no known brain mechanisms that can explain the discovery processes. Collecting many examples does something quite different.

I am groping towards an alternative account within the theory of evolved construction kits. Those construction kits seem to have produced a type of mathematical reasoning that is very different from a Turing machine, because the basic states and operations are different, e.g. because of the forms of representation and reasoning in brains required for making the ancient geometrical discoveries.

I suspect no variant of current models of neural mechanisms will suffice. It is possible that there is some important, currently unknown, feature of sub-neuronal chemical computation that will provide the answers.

After all chemical structures in genomes are somehow able to specify all the geometric structures of individual plants and animals, and also their models of control.

Triangle qualia
How do you work out the answer to this question (posed earlier):
If a vertex of a planar triangle is moved further from the opposite side of the triangle, and the sides remain straight, the angle at that vertex must get smaller. Why?

I suspect that many very bright children can "see" what the answer must be though they may struggle to articulate why that must be the answer. (Try it, without the use of pencil, paper, drawing in the sand, etc.!)

As David Hilbert pointed out, there are some analogies between such spatial/topological reasoning about triangles and reasoning in metamathematical proofs about derivability. E.g. discovering that it is impossible in a certain formal system to derive a particular formula from the axioms and rules. Such reasoning typically treats formulae and proofs as structures in a kind of discrete linear 'geometry'.

Anyhow that triangle example is just one among many examples I have been collecting, at the same time as trying to identify how biological evolution could have produced the required information processing machinery (implemented in physics and chemistry of course).
[Here's an example concerning rubber bands:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/rubber-bands.html]

It may take another century or so to work out all the details, especially if nobody else joins the project, in which case I'll require brain-in-vat support machinery to keep me going.

Perhaps a century is too optimistic. (Current arguments about an approaching "singularity", a point in time at which machine intelligence matches or overtakes human intelligence, all make false assumptions about how brains work: there is no evidence that their assumptions can explain the achievements or Archimedes, or squirrels, or elephants, or weaver birds, or this toddler apparently forming and testing conjectures in 3D topology.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#pencil)

A few more examples of different sorts:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html
The Triangle Sum Theorem (old and new proofs)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html
Hidden Depths of Triangle Qualia
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/shirt.html

Several kinds of gap in our knowledge

E.g. we don't understand:

(a) Forms of intelligence used by humans (young and old) and other animals
(including individual variations within species, and during development)
(b) What mechanisms make those forms of intelligence possible
(including forms of representation and manipulation of information)
(c) How those forms of evolution evolved, and
(d) How they develop in individuals
(e.g. linguistic competences and topological and geometrical reasoning competences in humans).

The M-M hypothesis

Perhaps understanding more about intermediate stages in evolution of various kinds of natural intelligence will draw attention to gaps in our ideas about modelling and replication of natural competences.

Including intelligence of various known and conjectured intermediate life forms in the more or less distant past -- sometimes inferred indirectly from likely challenges faced by evolutionary ancestors.

Thinking in detail about new challenges and opportunities produced by early evolutionary developments may provide new ideas about evolved solutions.

==

What Can We Learn From Conjectured Evolutionary Transitions?

Figure 1: Evo-transitions

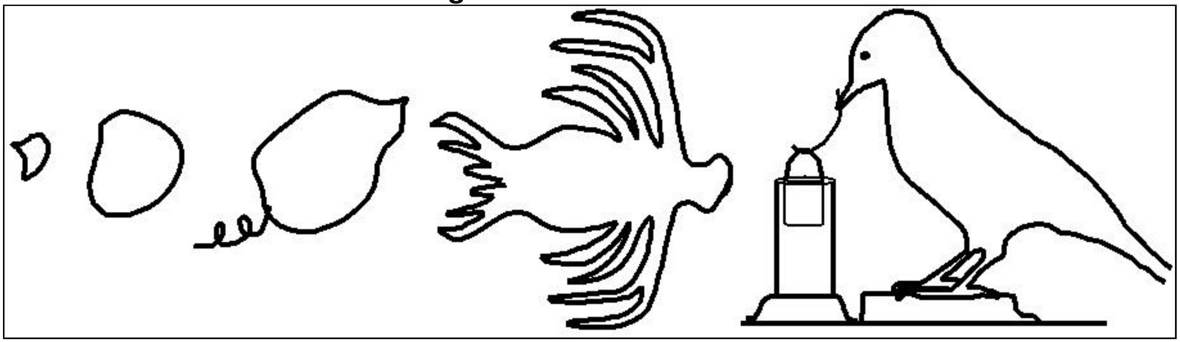

An example of crow intelligence (Betty reported in 2002):
https://www.youtube.com/watch?v=OYZnsO2ZgWo
For more information about Betty see Oxford Ecology Laboratory
Online videos show that she made hooks out of straight pieces of wire found nearby, in at least four different ways.

==

Examples:

- What new forms of intelligence were required by the transition from water-based to land-based-life, or from one kind of terrain to another, or when a new type of predator or prey turned up, or when articulated body parts became independently controllable, or when remote-sensing mechanisms evolved, or when various kinds of flight developed, or when offspring were born or hatched unable to feed, or to move, themselves, or when mating required cooperation?

- What new forms of intelligence were required for the ability to notice that some things are not merely supported or contradicted by evidence, but can be known to be possible or known to be impossible or known to be necessarily true or false, features that cannot be derived from empirical observation.

Problem: educational gaps in our culture

Part of the problem is that some of the hardest questions are not widely understood by researchers, because formulating the questions accurately requires concepts from epistemology, and metaphysics which hardly anyone learns at school, nowadays.

E.g. the distinction between empirical and non-empirical knowledge, and the closely related, but different, distinction between contingent and non-contingent (necessary) truths and falsehoods, both discussed by Immanuel Kant in 1781.

Necessity, impossibility, and related concepts have nothing to do with statistical evidence or probabilities.

Those concepts are often assumed to be based on "possible-world semantics" (e.g. what is true in this world and in all possible alternative worlds).

But for a child or non-human animal, or someone making a topological discovery, impossibility has nothing to do with alternative complete worlds.

Only possible alternatives to a localised portion of this world are relevant.
A more detailed, but still incomplete, analysis was presented in my 1962 DPhil Thesis, defending Kant.
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-1962

"If a problem is too hard to solve,
try to find a related harder one"
(I can't now recall where I learnt that.)
===

Where do various forms of intelligence fit in the scheme of things?

LEVELS OF STRUCTURE IN OUR PART OF THE UNIVERSE

Figure 2: A subset of branching metaphysical layers

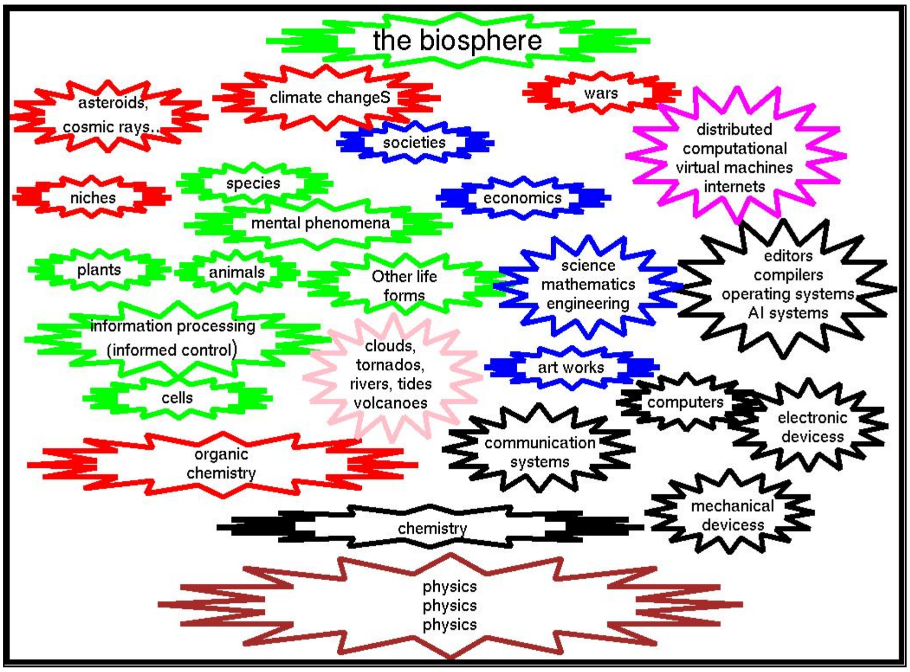

===

If we look beyond our planet we find far more structures of various kinds on various scales: with new discoveries being added constantly.

Those are structures found by physicists/cosmologists

E.g. see Max Tegmark's book:
Our mathematical universe, my quest for the ultimate nature of reality

What features make physical stuff able to implement and bring into existence implementations of minds and their properties?

All the above summarised structures on our planet must somehow have been produced by the structures studied by physicists.

Any adequate fundamental physical theory therefore needs to be capable of playing a deep role in explanations of the origins and workings of all the structures on our planet (and perhaps other more complex structures in other parts of the universe, unknown to us).

One of the things we have learnt from developments in computer science and engineering in the last 70 years or so is that there are different sorts of deep roles that a single powerful explanatory mechanism can have, illustrated by the right side of Figure 2.

For example there is now a huge variety of uses of computers, differing enormously in physical scale, in type of function, in kinds of applications, in functions served: and all of these depend on a common substratum that can take different but mathematically equivalent forms, including large arrays of bistable switches that can be made to change their on/off patterns under the control of other switches.

==

WADDINGTON'S EPIGENETIC LANDSCAPE

Figure 3: Waddington

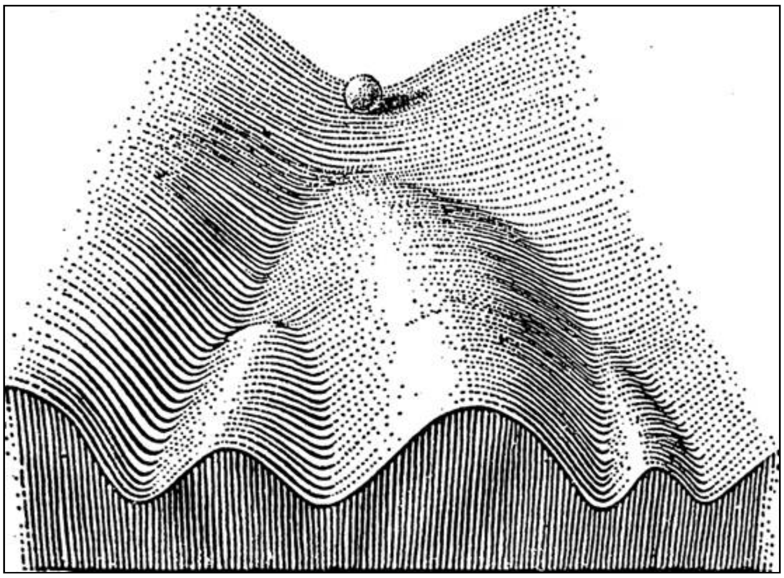

==

AN ALTERNATIVE VIEW OF EPIGENESIS

Figure 4: Chappell-Sloman Epigenesis
Multiple developmental trajectories from one genome,
with new branches triggered at various stages of development:
using landscape changes produced at different stages
during development (new barriers, routes, tunnels...)

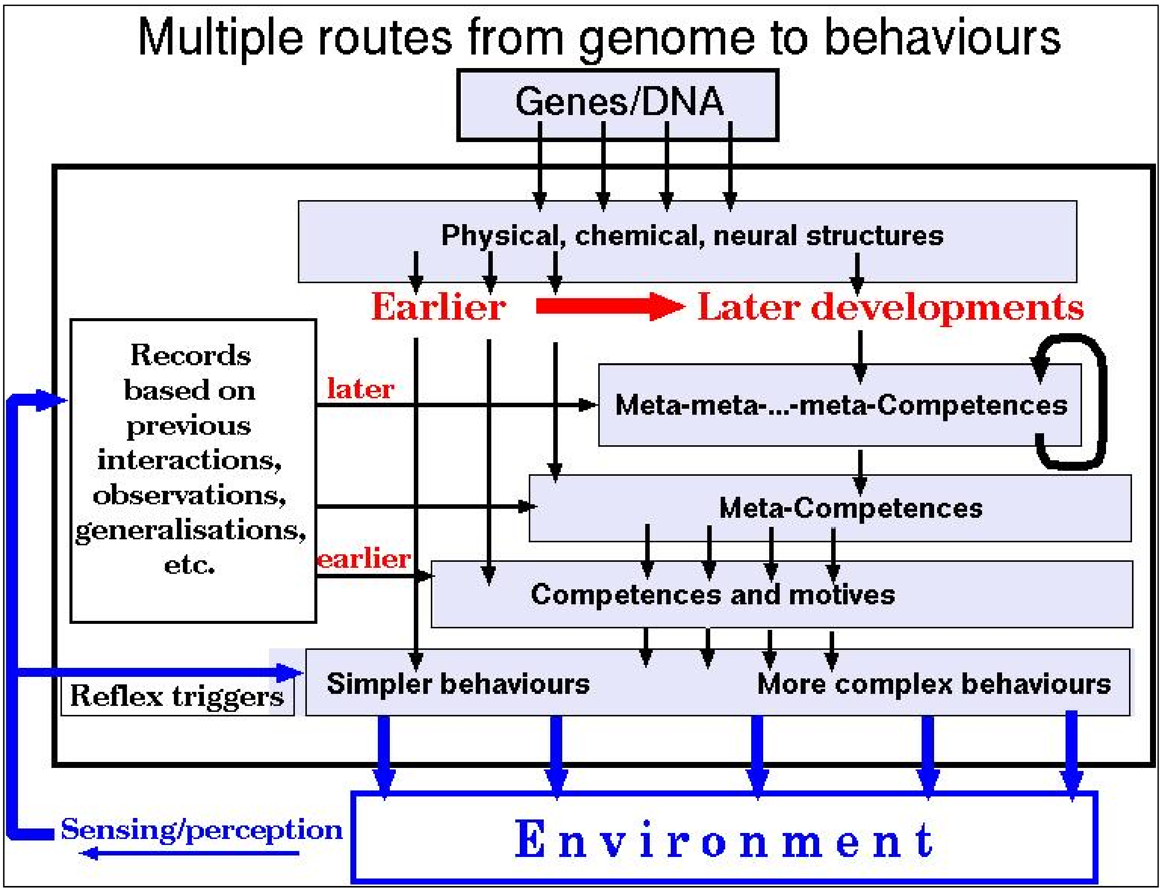

Figure adapted from: Jackie Chappell and Aaron Sloman, 2007,
"Natural and artificial meta-configured altricial information-processing systems,"
International Journal of Unconventional Computing, 3, 3, pp. 211--239,
http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717

==

Can we extend some of those ideas about epigenesis in organisms to epigenesis in the universe as a whole?

For now, let's consider this only in relation to evolution on one planet,
using ideas developed in the Meta-Morphogenesis project
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html
especially the work on fundamental and derived construction kits used by natural selection.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
(still extending ideas published in Sloman(2013), and Sloman(2017)

More precisely, are there parallels between the multiple routes from a genome to behaviours of organisms in the same species and multiple routes from fundamental physical properties of the universe to products of evolution in different locations or different sorts of different environments?

Figure 5: Cosmic epigenesis (a)
A step towards generalising Figure 4

Multiple routes from baby universe to layered products

Proto-Stuff in proto-Space-Time
(The Fundamental Construction-kit FCK)

Physical/chemical/cosmological processes and their
multi-scale products, including
Derived Construction Kits (DCKs)
e.g. chemical construction kits

In at least one small part of the universe (perhaps in many parts?):

Physical chemistry leads to organic chemistries

A sketch of epigenesis on a universal scale (a)

Figure 6: Cosmic Epigenesis (b)

In at least one small part of the universe (perhaps in many parts?):

Physical chemistry leads to organic chemistries → **Carbon-based / Silicon-based? Others?**

Supporting:
Increasingly complex and varied construction kits, used for increasingly complex and varied mechanisms and life forms

Using: **Increasingly complex physical structures and mechanisms making and using new materials with new functions.**

Producing: expanding demands for energy stores, available when needed.

More complex and varied physical mechanisms:

- **Self-assembled body-parts, sensors, effectors, information stores, nutrient allocation, infection control, waste disposal, damage repair, reproductive mechanisms, ...**
- **Increasing (space and time) gaps between acquisition and use of materials and information by individual organisms (varied forms of offline intelligence)**
- **Increasing use of other organisms for materials, energy, protection, many forms of exploitation and collaboration**
- **More subtle, abstract, and context sensitive genomic information, including delaying gene expression while information about current context is acquired (human language an obvious case)**
- **Especially structural rather than numeric information (meccano, lego...)**
- **Increasingly powerful information-based control mechanisms: initially physical control mechanisms, then later physically implemented (multi-layer) virtual machinery, using delayed development to ensure tailoring to current context.**
- **Followed later by metacognitive mechanisms allowing self debugging, reflective analysis, generalisation, understanding why things fail or succeed (using more layers of metacognition)**
- **Later on other-directed metacognition used in competing, teaching, collaborating, and culture-formation.**

More and more Meta-meta-...-meta-Competences
used by great ancient mathematicians

A sketch of epigenesis on a universal scale (b)

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, At that stage options for life forms in the short term were very limited.

(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.
These rapidly expanded the "reachable" options for new designs with new functionality.
Compare the advances in types of computational machinery made possible as new physical construction kits, new low level programs, new programming languages, new operating systems, new file systems, new compilers, new interpreters, new networking technology, new forms of virtual machinery, became available e.g between 1950 and now.

The vast majority of computational machines and applications, concrete and virtual, which are now actually in use and likely to appear in the next few years were in a deep sense completely unreachable in 1950 without first going through many intermediate stages of discovery and development.

Figure 7: The Fundamental Cosmic Construction Kit

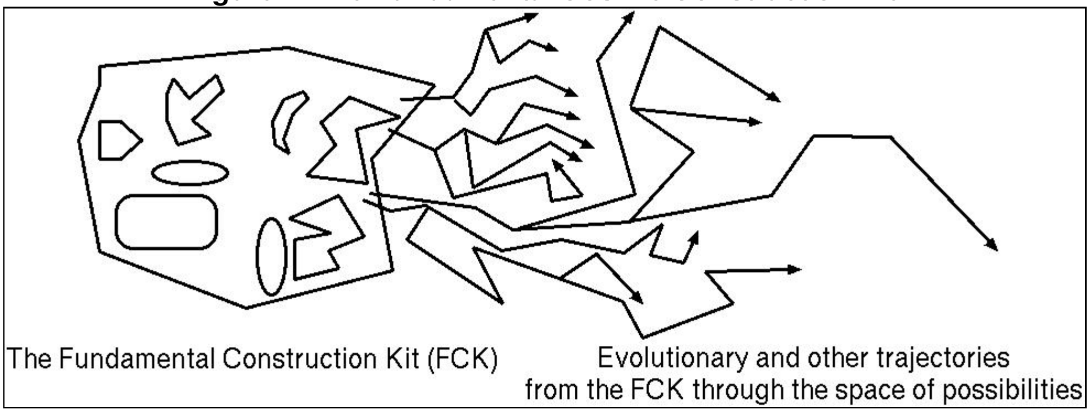

Figure 8: Adding Derived Construction Kits

This shows only one new derived construction kit. The theory of evolved construction kits postulates the use of many layers, evolved at different times, including concrete, abstract and hybrid construction kits required for increasingly complex forms of information processing.

The theory of construction kits is still under development, though partial progress reports have been made available, in a Springer collection published in 2017 (based on the theory in 2016), and an ongoing progress report on the theory of evolved construction kits: Sloman(2017)

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.

In part this results from the use of abstraction in many of evolution's discoveries: since abstracting from specific details forms an essential part of a design for a control mechanism for something that can grow, or change its shape. This is discussed in more detail in Sloman(2017). Abstraction, based on use of mathematical structures, seems to be at the heart of evolution's ability to produce a huge variety of effective designs. This is why I sometimes refer to evolution as a "Blind Mathematician". (Compare Dawkins' "Blind Watchmaker" analogy.)

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

also including topological reasoning about effects and limitations

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.

TO BE COMPLETED

REFERENCES AND LINKS

- Jeremy Bernstein, Profiles A.I. (Marvin Minsky interview) December 14, 1981 *New Yorker* http://www.newyorker.com/magazine/1981/12/14/a-i
- H. Gelernter, 1964, Realization of a geometry-theorem proving machine, reprinted in *Computers and Thought,* Eds. Edward A. Feigenbaum and Julian Feldman, McGraw-Hill, New York, pp. 134-152,
  http://dl.acm.org/citation.cfm?id=216408.216418
- Immanuel Kant, 1781, *Critique of Pure Reason*
- Aaron Sloman, *The Computer Revolution in Philosophy, Philosophy, Science and Models of Mind,* Harvester Press (and Humanities Press), 1978, Revised 2015, Hassocks, Sussex,
  http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp
- Aaron Sloman, 2013, 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, Elsevier, Amsterdam, pp. 849-856,
  http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d
- Aaron Sloman, 2017,
  Construction Kits for Biological Evolution, in *The Incomputable: Journeys Beyond the Turing Barrier,* Eds. S. Barry Cooper and Mariya I. Soskova,
  NOTE: the published versions are mere snapshots of a continuing process of development.
  http://www.springer.com/gb/book/9783319436678
  Ideas still being developed at:
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
- M. Tarsitano, 2006, Route selection by a jumping spider (Portia labiata) during the locomotory phase of a detour, *Animal Behaviour,* 72, Issue 6, pp. 1437--1442,
  http://dx.doi.org/10.1016/j.anbehav.2006.05.007

Installed: 4 Jul 2017
Last updated: 10 Jul 2017; 16 Jul 2017; 29 Jul 2017; 5 Aug 2017;

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