School of Computer Science

The Meta-Morphogenesis (M-M) Project
Alternative title:
The Self-Informing Universe

Was Alan Turing thinking about this before he died, two years after publishing his 1952 paper on Morphogenesis? [*]

Aaron Sloman
http://www.cs.bham.ac.uk/~axs
Short link to this site: http://goo.gl/9eN8Ks

A Protoplanetary Dust Cloud?

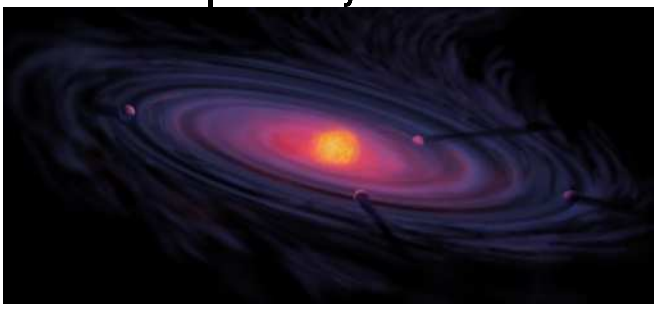

[NASA artist's impression of a protoplanetary disk, from WikiMedia]

KEY QUESTION
How can a cloud of dust give birth to a planet
full of living things as diverse as life on Earth?
Including great ancient mathematicians, e.g. Euclid, Archimedes, Zeno, etc.
and highly intelligent nonhumans, e.g. squirrels, elephants, crows, ...

JUMP TO CONTENTS LIST

BACKGROUND NOTES
This document has been under construction, and revised several times, since the project began in 2011, triggered by Alan Turing's centenary, especially an invitation to comment on Turing's 1952 morphogenesis paper, as explained below.

Some new lines of thought began when I started thinking about metamorphosis in insects, in June/July 2023, as reported in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/metamorphosis.html
That is now (December 2023) my most complex online document, with links to information from a wide variety of disciplines and researchers.

New sub-projects and subsidiary papers and videos are spawned from time to time, referenced below and in related documents. This document also builds on older, related work going back to my first encounter with AI while a young philosopher around 1969, inspired by Max Clowes. A personal tribute to Max Clowes can be found here:
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html
(Also PDF)

Note: updated 23 Dec 2023
Earlier changes
Major new developments began around October 2020, related to the problems of explaining how processes that occur in a hatching (vertebrate) egg can produce cognitive competences ready for use shortly after hatching, i.e. with no need to train neural networks first. The ideas continued evolving during 2021 and 2022, including possible challenges for fundamental physics, and claims that controlling processes inside the eggs are related in very complex and unobvious ways to stages in the evolution of the species.

A messy overview document (produced in October 2022 -- much revised since then, and still under revision) is available here:
https://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-devo.html
including some very complex (over-complex?) diagrams attempting to convey some of the complex transitions during evolution of egg-laying vertebrate species and how they are related to increasingly complex biochemical transitions during hatching processes inside eggs.

20 Jul 2023: Another example: Starfish embryos
Some time ago, Leo Caves kindly sent me information about this video, which I had previously mislaid:

https://youtu.be/bki2kI8aQvg
showing crystallogenesis and large-scale wave motion of starfish embryos.

Note added: 17 Oct 2020, updated ... 28 Oct 2022 ... 19 Jul 2023
The above new developments in this project were partly triggered by thinking about what Schrödinger did and did not say about life and evolution in (1944), and what he could have said.

I shared some half-baked thoughts about that with Anthony Leggett (theoretical physicist at UIUC). His combination of scepticism and encouragement triggered important modifications and extensions of the proposed theory.

Several changing talks about these ideas were given during 2020, 2021, 2022 and 2023, presenting a sequence of modifications in the ideas, especially new ideas about how the extremely complex, multi-layered processes of development in a vertebrate foetus inside an egg, correspond to multi-layered evolutionary processes extending the in-egg assembly mechanisms producing a new animal, that not only has a highly complex, highly functional physical body with intricately related internal components and also complex competences available shortly after hatching (including walking and foraging competences in some cases) that could not be based on learning processes after hatching.

For a messy, still evolving, overview of the new ideas presented in 2022 see the document referenced above.

These ideas are also extensions of Sloman(2020), on the evolution of multiple varieties of consciousness.

Those ideas focus on the observation that many newly hatched animals, developed inside eggshells, have significant forms of spatial intelligence that they could not have learnt after hatching and could not have acquired inside the egg by any of the currently proposed learning mechanisms, which require perceiving and acting on the environment, to acquire statistical evidence from which probabilities are derived, e.g. using neural networks. The developing embryo inside an egg obviously cannot do that.

An example showing some capabilities of newly hatched avocets is in this 35 second videoclip from a BBC Springwatch programme in June 2021:
https://www.cs.bham.ac.uk/research/projects/cogaff/movies/avocets/avocet-hatchlings.mp4

Such spatial intelligence available soon after hatching must be a product of *multi-layered* chemical assembly processes inside the egg. I don't believe any current theory of cognition, or chemical synthesis, explains how that can happen.

This challenge, to explain how eggs can produce both physiologically fully formed (but still immature) hatchlings with significant cognitive (spatial) competences available soon after hatching, is a recent extension to the (Meta-Morphogenesis) project.

An invited talk about this, presenting earlier versions of the ideas about hatching, was given via zoom at Tuebingen university in mid 2021, recorded here:
https://www.youtube.com/channel/UCaG1Q8TEuLN5OJZXaTL28PQ
(Screen sharing failed, so it is a talking-head presentation).
The discussion following the talk is recorded here:
https://www.cs.bham.ac.uk/research/projects/cogaff/movies/tub/as-discuss.mp4

Later Zoom Talk on 21st June 2022

A later version of the theory was presented in an invited talk plus discussion (over two and a half hours) on 21st June 2022 hosted by a biomedical research group in Singapore. A recording of the talk and discussion is available here:
https://www.a-star.edu.sg/bii/highlights
In case of problems try this instead:
https://drive.google.com/file/d/1hyjam-wBEx7zIir_8TNRbYF14vddsNZa/edit
or this
https://zoom.us/j/91429797476?pwd=TW5sTElsUG0zSkhrbm5ITGkzbkZVZz09

Talk at Tokyo Conference January 2023

A more recent, shorter (zoom) presentation was given at a Tokyo conference in January 2023. Unfortunately the recordings of talks at that conference are not freely available. The talk was based on an earlier version of this (still evolving) overview:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-devo.html

Leggett Comments, January 2023
Theoretical physicist, Anthony Leggett, was invited to comment on my Tokyo talk. Unfortunately, he was not able to connect at the time of the event, so he provided written comments, responding to written notes on my talk topic, sent to him before the conference. His comments, read out by the organiser during the workshop, are available here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/tony-leggett-talk-notes.txt

N.B.This is all still work in progress, and my ideas have developed since then, including thoughts about metamorphosis in insects, mentioned above.

Note added: 22 Jan 2020
I have found this presentation (with transcript), by Neil Gershenfeld (MIT), with several points of overlap:
Morphogenesis for the Design of Design, (Edge Talk, 2019),
https://www.edge.org/conversation/neil_gershenfeld-morphogenesis-for-the-design-of-design

---

The M-M project
Although partly based on work done within the framework of the CogAff (Cognition and Affect) project, the Meta-Morphogenesis project began with a new name and a new focus late in 2011 -- initially as an invited chapter published in 2013 in part IV of the award winning Elsevier Turing Centenary volume.

> Aaron Sloman, 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, 2013, pp. 849-856,
> http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d
> The contents of the book are listed here, with additional information:
> http://www.cs.bham.ac.uk/~axs/amtbook

That chapter, an attempt to understand why Alan Turing had worked so hard on chemistry shortly before his death, offered the conjecture that if he had lived a few more decades he would have extended our ideas of computation (information processing) by showing how chemistry-based information processing, with its mixture of discrete and continuous processing, plays an essential role in the mechanisms of biological evolution and the development of its products. I proposed the M-M project to continue his pioneering work. It was preceded by two scene-setting chapters in the same book,

> Virtual Machinery and Evolution of Mind (Part 1), pp. 97--102,
> http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106a

> Virtual Machinery and Evolution of Mind (Part 2), pp. 574--580,
> http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106b

Since then, the M-M project (like this document) has been revised, expanded and split into sub-projects several times. An important sub-theme discusses the essential roles of "Construction Kits" (fundamental and derived, concrete and abstract) used by and (mostly) created by biological evolution. An early version of the Construction-Kits paper, frozen in 2016, was published by

Springer in 2017: https://doi.org/10.1007/978-3-319-43669-2_14, as part of a book, edited by S. Barry Cooper and Mariya I. Soskova:
https://www.springer.com/gb/book/9783319436678

Like everything I write, the published version was soon out of date. Since then the online version of the construction-kits paper has continued growing, and spawning sub-papers, here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
(Also pdf.)
Notes for a zoom presentation in Tuebingen on 2nd June 2021 can be found here (updated after the presentation):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/unsolved-talk.html

Different aspects of the Meta-Morphogenesis project
There are many ways of viewing the project. One view emphasises biological evolution's increasingly sophisticated and varied uses of information -- semantic, usable, information of the kind often referenced by Jane Austen in *Pride and Prejudice* rather than the syntactic variety Shannon identified over a century later -- confusing a whole generation of scientists and philosophers (though he wasn't confused). That thread is discussed in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html.

Another strand is evolution's increasingly sophisticated use of varieties of mathematics -- not the logic-based algebraic kind of mathematics that now dominates much thinking about about the nature of mathematics, but the ancient uses of mathematical understanding of spatial (geometrical and topological) structures and processes and their impossibilities and necessities in the discoveries of Archimedes, Euclid, Zeno and many others (emphasized by Immanuel Kant in his philosophy of mathematics, nowadays mistakenly thought to have been refuted by Einstein and Eddington). An example of such an ancient discovery is the *neusis construction*, known since before publication of Euclid's Elements

One aspect of evolution's use of mathematics is its repeated discovery of increasingly powerful re-usable mathematical abstractions shown by (among other things mentioned below) its reliance on growing layers of compositionality
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/compositionality.html
(also pdf),
manifested in, among other things, its use of meta-configured genomes (MCGs) (using multi-layered, staggered expression of increasingly powerful genetic/epigenetic abstractions), outlined (in collaboration with Jackie Chappell) here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html
(also pdf).

[Without such a cross-generational boot-strapping mechanism, human languages would have had to be taught to humans by alien experts in those languages -- of which there are now several thousand, including many sign languages.]

Another role for mathematics is in the spatial intelligence of many species including squirrels, crows, elephants, orangutans, and pre-verbal human toddlers, like this toddler topologist (aged about 18 months):
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/small-pencil-vid.webm whose spatial understanding I suspect is unmatched by any robot on the planet.

That’s true also of spatial understanding in many other intelligent species mentioned above, although there are robots with specialised competences no animal can match, e.g. used on car assembly lines. But they know not what they do or why it works!

New links to theories and models in developmental biology
In 2019 Peter Tino pointed out that the ideas in this project seem to be related to mechanisms observed some time ago (to the great surprise of biologists) in gene transcription - vastly enriching the ways in which proteins based on genome codes can be produced, substantially extending initial ideas about genome-triplets of base pairs specifying aminoacid enabling protein synthesis. (I don’t have expert knowledge in that area but I have been trying to document some of the connections in updates to the M-M project.)

> Example quote from the Smithsonian Institute:
> http://humanorigins.si.edu/dna-and-evolution
> "An older theory of genetics maintains the principle of "one gene, one protein". However, modern genetics has discovered that oftentimes, proteins are determined by the coordinated activities of several genes."
>
> The meta-configured genome theory shows (in principle, but not in detail) how the more complex mechanisms can contribute to the importance of compositionality in biological evolution.
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/compositionality.html

After an encounter with a new pathogen, the adaptive immune system often "remembers" the pathogen, allowing for a faster response if the pathogen ever attacks again.

> One of many web tutorials on immune systems.
> (Khan Academy)

This is *chemical*, not *neural*, remembering. We suggest that, in a similar way, a shared *meta-configured* genome implemented chemically can produce dramatically different effects in different physical terrain requiring different forms of behaviour, or in different cultures, in the same species, though using far more complex mechanisms than antibodies.

Perhaps the most striking example is the ability of the multi-layer human genome to produce competences in thousands of very different languages, that differ at all levels from the basic sounds used (or signs, in the case of sign-languages), up to modes of composition of phrases, sentences and beyond (e.g. However..., Therefore..., ... although...).

So, instead of all learning being based on abstractions and probabilistic generalisations from statistical data collected by individuals (as in currently proposed neural net models), we suggest that chemistry-based ’compositional’ processes and mechanisms are involved in development of many forms of information processing, including linguistic competences and increasingly sophisticated types of proto-mathematical spatial reasoning competences used by ancient geometers (Archimedes, Zeno, etc.) -- and their precursors -- that don’t seem to be explicable using familiar logic-based, rule-based or neural net based AI mechanisms.

The development of immune systems that such mechanisms make possible is not done by learning in the individual, but by evolutionary learning in the species, possibly over many generations, and across multiple layers of species evolution.

Perhaps many cognitive functions in intelligent animals have a similar long term aetiology -- based on chemically implemented multi-layer *species* learning, tracking physical and cultural changes in the environment, rather than *individual* learning.

As Hume pointed out, no "ought" can be derived from "is", yet many animals have mechanisms generating motives, preferences, attitudes, emotions, and other affective states and processes, all products of biological **control** processes of different sorts. The investigation of varieties of control, their uses and the mechanisms making them possible is the core of the Cognition and Affect project summarised in this (unfinished) overview
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/emotions-affect.html (also pdf)

Earlier work, in the CogAff project (papers and presentations) is available here:
http://www.cs.bham.ac.uk/research/projects/cogaff/
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/

A key feature of the earlier work (since my 1978 book, The Computer Revolution in Philosophy) is investigation of information processing architectures, i.e. attempting to understand what sorts of information processing mechanisms are required (e.g. to make a human-like mind possible), how they have to be combined to form a functioning mind, how such mechanisms evolved, and how they can be implemented -- in particular whether they can all be implemented using digital computers, or whether some other type of mechanism is required for some sub-systems, e.g. chemical mechanisms (known to be central to control of reproductive processes, as discussed in Schrödinger (1944)). Organisms (like many complex human-designed systems) require concurrently active layers of physical and virtual machinery. Different layers were produced at different stages of evolution, with higher layers implemented in but not reducible to lower layers. Some discussion of the issues is in this paper on Virtual Machine Functionalism (also pdf):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html

Compositionality
Another new sub-theme was triggered in September 2018, namely, evolution's use of compositionality, assumed implicitly in earlier work but now discussed explicitly here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html
Also PDF
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.pdf.

Meta-Morphogenesis and Metaphysical Grounding
I have been interacting with the Birmingham FramePhys project
https://www.birmingham.ac.uk/schools/ptr/departments/philosophy/research/projects/framephys/index.aspx
led by Alastair Wilson. Some notes following a meeting of the project can be found here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/emergent-physics.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/emergent-physics.pdf

Earlier notes on a "Metaphysical Grounding" sub-theme
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-framephys.html were inspired by discussions in Jan 2019 at the FramePhys+Gothenburg Conference on Metaphysical Explanation

in Science
https://www.birmingham.ac.uk/schools/ptr/departments/philosophy/events/2019/framephys-conference-bham-1001.aspx
organised by the Birmingham https://twitter.com/FraMEPhys project and the Gothenburg Metaphysical Explanation project. These notes Sloman(2019), summarise some of the links between the Meta-Morphogenesis project and Metaphysics, pointing out deep connections between biological evolution's discoveries, including its creative use of mathematics, and the metaphysical status of mathematical truths.

Other sub-themes include deep limitations of deep learning, Jane Austen *vs* Claude Shannon on the nature of information, the Meta-Configured genome, "architecture-based" *vs* "reward-based" motivation, why Immanuel Kant's philosophy of mathematics was not refuted in the 20th century by Einstein and Eddington, transitions in the space of information-processing architectures, the basis of *alethic modalities*, and several more, all pursued in papers, talks and videos in the Cognition and Affect (CogAff) web site, the implications for cognition of Erwin Schrödinger's discussions of quantum mechanisms and chemistry in 1944, the creativity of biological evolution.

What's the connection with Turing?
Evidence (admittedly inconclusive) for Turing's interest in aspects of this project concerning non-digital (e.g. chemistry-based), forms of computation, can be seen in various publications, especially his last major publication Turing(1952) (usefully summarised for non-mathematicians by Philip Ball in Ball, 2015). There is also an unexplained sentence and the single sentence mentioning chemistry in (Turing, 1950): "In the nervous system chemical phenomena are at least as important as electrical".

A clue, of which I was unaware until 2018, is Turing's very brief discussion of a distinction between mathematical intuition and mathematical ingenuity, in his 1938 PhD thesis, where he claimed, without explanation, that only the mathematical ingenuity, not mathematical intuition, can be implemented in (digital) computing machinery. He said nothing there about how intuition might be implemented in brains and why not in computers, though there is an unexplained remark about the importance of chemistry in brains, in his Mind 1950 paper referenced below. He may have been thinking of the fact that chemical mechanisms include a mixture of discrete (e.g. catalytic) processes, and continuous processes (as molecules, or parts of molecules, move through space, become twisted, etc.

I have tried to summarise and discuss the evidence from the brief remarks in his thesis here (added late 2018):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html (and pdf)

Less direct evidence for relevance to Turing's thought is provided by asking: Why did Turing produce a major piece of work on chemical morphogenesis, Turing(1952), two years before he died, and what might he have done if he had lived several decades after that?

My tentative answer is that if Turing had lived a few decades longer, he would have explored the potential for chemistry-based information processing, using a mixture of discrete and continuous machinery with features like those highlighted as crucial for biological reproduction and evolution by Erwin Schrödinger: (1944). It is also possible that before he had died he had read Kenneth Craik's little book Craik(1943), which raises questions about the role of chemistry in brains. But as far as I know there is no evidence that the two ever met. Craik also died young, in a road accident, in 1945. His work had strongly influenced the founders of the Ratio club, of which Turing became a member later. Even if Turing and Craik never met, it is very likely that Turing learnt about Craik's work, e.g. through the Ratio club.

What's information?

A key assumption here is that the *core* concept of information is not concerned with *communication*, *storage* and *retrieval* of information, but with various *uses* of information, e.g. in taking decisions, making and using plans, controlling behaviour at various levels of abstraction, formulating questions, recording acquired knowledge, deriving consequences, formulating explanations and predictions, reflection on implications of recently recorded information, and many more. This point is expanded below, where I contrast Shannon's concept of information with the concept used a century earlier in Jane Austen's novels.

CHALLENGES:

This project challenges the thesis that digital computational processes can model everything in this universe with any desired level of precision. This rejected thesis, which I previously defended, e.g. in Sloman (1978), is implicitly or explicitly accepted by many distinguished scientists, including physicists who believe that in principle every detail of our universe can be modelled perfectly on a sufficiently large digital computer (or Turing machine). Examples can be found in
https://en.wikipedia.org/wiki/Digital_physics
(However, parts of that article are imprecise imprecise and it leaves out important details.)

A recent presentation of the claim that the physical universe can be completely modelled on a sufficiently large digital computer, is by the distinguished physicist **David Deutsch**, who has deep philosophical interests. Below is a quotation from his book,
The Beginning of Infinity: Explanations That Transform the World, published by Allen Lane & Penguin Books, 2011, page 458:

> "In regard to understanding the physical world, we are in much the same position as Eratosthenes was in regard to the Earth: he could measure it remarkably accurately, and he knew a great deal about certain aspects of it - immensely more than his ancestors had known only a few centuries before. He must have known about such things as seasons in regions of the Earth about which he had no evidence. But he also knew that most of what was out there was far beyond his theoretical knowledge as well as his physical reach.
>
> We cannot yet measure the universe as accurately as Eratosthenes measured the Earth. And we, too, know how ignorant we are. For instance, we know from universality that AI is attainable by writing computer programs, but we have no idea how to write (or evolve) the right one. We do not know what qualia are or how creativity works, despite having working examples of qualia and creativity inside all of us. We learned the genetic code decades ago, but have no idea why it has the reach that it has. We know that both of the deepest prevailing theories in physics must be false. We know that people are of fundamental significance, but we do not know whether we are among those people: we may fail, or give up, and intelligences originating elsewhere in the universe may be the beginning of infinity. And so on for all the problems I have mentioned and many more."

Note added 6 Jan 2022
His more recent work on what he calls "Constructor Theory" abandons those ideas. E.g. this interview provides a summary:
https://www.youtube.com/watch?v=d8KCXgzqhSQ
David Deutsch on 'Constructor Theory' 8 Dec 2018

But do we really know that (human-like, or squirrel-like) AI is attainable on digital computers? We do know that human brains include vast numbers of molecular (e.g. chemical) processes combining discreteness (e.g. catalytic reactions), continuity (folding, twisting, moving closer or further), enormous parallelism, and (quantum) randomness. Could any single, sequential Turing machine, with its discrete tape and discrete collection of symbols and discrete, deterministic, machine table, model such a brain precisely? The investigations reported in this project suggest not. John von Neumann raised related doubts in 1958. Compare Newport(2015). Even Alan Turing, in his 1938 PhD thesis, mentioned above, distinguished roles for *intuition* and *ingenuity* in mathematics, suggesting that computers can provide only the latter. What brain mechanisms are needed to provide mathematical intuition, and whether they can be replicated artificially is an open question discussed in various parts of this web site.

My own interest in these questions began with my attempt to defend Immanuel Kant's philosophy of mathematics in my 1962 DPhil thesis, written before I knew anything about computers or AI.

Michal Forišek wrote, in
https://www.quora.com/What-are-the-philosophical-implications-of-computational-universality-Turing-completeness

> .... "if you actually want to build something that will compute in a different way, you have to start by looking at physics: find something with a behavior fundamentally different from the sequential deterministic stuff used to construct current computers, and then come up with a way to use that new thing to do useful computations for you."

I suspect that shortly before he died, Alan Turing was beginning, in his 1952 paper, to explore an alternative based on the combination of continuity, discreteness, and unsynchronised parallelism found in chemical reactions, previously noticed by Schrödinger in *What is life?* (1944). This project is my pathetic attempt to guess at some of what Turing might have achieved, had he lived a few decades longer. I think Turing was not committed to the strongest form of what is now (misleadingly) called the "generalised" Church-Turing thesis. Early evidence for this disagreement is in his 1938 distinction between *intuition* and *ingenuity*, summarised and discussed in the paper referenced above.

NEWS JANUARY 2020
Added link to Neil Gershenfeld, Morphogenesis for the Design of Design, Edge Talk, 2019: which raises important relevant issues, from an engineering point of view.

NEWS OCTOBER 2018
Since a paper on compositionality was accepted for a workshop in Sept 2018, I have realised that there are deep and varied mathematically rich forms of compositionality in evolution, including compositionality in individual designs, in epigenetic processes, in biological niches, in ecosystems, as well as processes in individual minds, where compositionality was first identified, e.g. linguistic compositionality. This new theme is being developed in a separate paper:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html, also PDF

CONTENTS

Updated: 6 Jan 2022 Extended references to Deutsch on Constructor theory. My own work on fundamental and derived construction kits, continues to grow.
22 Nov 2018 (new short abstract & minor reorganisation); 4 Mar 2019 (clarified abstract); 19 Jun 2019 (restructured)
5 Oct 2018 (added compositionality); 4 Nov 2018 (Turing on intuition);
Reorganised Mar, Apr, Jun, July 2018; 12 Aug 2018; 28 Sep 2018;
13 Jan 2017; 3 Jun 2017; 28 Jul 2017; 27 Dec 2017; 15 Jan 2018;
27 Oct 2016; 4 Nov 2016; 5 Dec 2016; 11 Dec 2016;
19 Dec 2014; 16 Mar 2015; 9 May 2016; 18 Jun 2016); 12 Aug 2016;
25 Aug 2014
Some of the contents have been moved to a file containing references:
Publications and references related to the Meta-Morphogenesis Project
Including a section on links to related projects (added 19 Dec 2014).
See also:
Background Notes (above)
UPDATES (A partial list) at the end of this file

CONTENTS LIST ABOVE

New summary (revised 30 Nov 2018; 4 Mar 2019)
(Revision of summary still in progress)

This project is about the evolution of increasingly complex and varied information-processing mechanisms operating on increasingly complex and varied forms of information in increasingly complex and varied living organisms, as a result of still unknown processes that began in a lifeless universe billions of years ago. What features of that lifeless universe made possible a collection of increasingly complex boot-strapping processes that produced not only a huge variety of physical forms of life, but also the still largely uncharted variety of types of information, and information-using mechanisms that made those life forms possible?

[Jane Austen's concept of usable information is relevant here, as explained below, not Shannon's syntactic concept.]

How can the huge diversity of life forms, with increasingly complex and varied combinations of physical forms, behavioural capabilities, habitats, types of information and information-processing machinery (including virtual rather than physical machinery) be encompassed in an integrated theory, which includes both descriptions and explanations of the details and explanations of how the details are constantly being extended, over time in the life of each individual, and over generations within (branching) species?

Many (most?) philosophies of learning and science assume there's a fixed type of world to be understood by a learner or a community of scientists, using fixed, general-purpose learning and discovery mechanisms, e.g. discovery of sometimes co-occurring features, and derivation of probabilities of co-occurrence.

In contrast, biological evolution, in collaboration with its products, is metaphysically creative, and constantly extends the mathematical diversity and complexity of the world to be understood, including constantly extending the types of learners with new powers, whose actions and, the products of their actions, contribute to extending what needs to be learnt by new members of the species.

At first this happens mainly by change of environment and/or evolution of new species, leading to new genome/environment interactions expressed in individual behaviours that differ across generations.

In more advanced (late-evolved) learners, a multi-stage/multi-layered genome can extend the powers of each learner in different ways at different stages during epigenesis -- producing creative, self-modifying, developmental trajectories influenced by both the (possibly unique, or dynamically changing) local environment and also the epigenetic stage reached by the individual -- just as ancient genomes produced chemistry-based construction-kits for building different body parts at different stages of development, and different control mechanisms for producing different internal or external processes at different times.

(Compare the diversity of types of information and types of control involved in, breathing, sucking, grasping, kicking, crawling, walking, noticing an object, moving to an object, touching an object, grasping an object, manipulating an object in various ways, dismantling an object, causing two objects to interact, preventing interaction, combining objects into larger objects, ... and what about microbe actions?)

Layered/staggered gene expression allows the same (or approximately the same) genome to produce very different developmental trajectories in different individuals or in different environments -- with resulting differences in physical development, behaviours, forms of learning, and individual actions on the environment.

So powers achieved by a learner at any stage in expression of a multi-layer genome, are influenced by a combination of
(a) general features of the environment,
(b) earlier learning powers, acquired during earlier stages of gene-expression, and products of resulting interactions with their environments
(c) new forms of gene-expression produced by interaction of later gene-expression with products of (a) and (b).

As a result, the gene-expression processes within members of a species can differ significantly across co-existing individuals, with slightly different environments, and also across generations, across geographical locations, and across cultures, changing dramatically under the influence of new technology produced by the species, as humans demonstrate so clearly. (Some minor differences may be due to differences in individual genomes. But substantial differences in cognitive abilities, motivation, and achievements between a five year old human two thousand years ago in ancient Egypt, a five year old living now in an isolated village with no electricity, telephones, or internet and a five year old growing up in a radio/TV/internet-enriched home in 2019 will mostly not be due to genetic differences, but to the materials available to be used by successive stages of gene expression in the first five years of life.

There were earlier "less creative" stages of the species where individual or group actions could not have such dramatic effects on gene expression -- e.g. ancient humans with relatively unchanging environments and cultures. For humans, especially those alive during and after the last few decades of the 20th century, the rate of production of new technology, with costs for users going down while powers and applications increase, and geographical availability accelerates, has produced unprecedented rates of change in what individuals can do, and thereby altered much of what they want to do.

Despite vast amounts of research, much is still unknown, including key features of the evolved biological mechanisms that underly many varieties of spatial reasoning in human and non-human species, including the mechanisms in human brains and minds, partly shared with other intelligent species, which, in humans, produced minds of ancient mathematicians whose discoveries laid foundations for much of what followed, using information processing mechanisms that (I claim) are not yet understood and not yet replicated in machines.

One of the key explanatory ideas, developed in collaboration with Jackie Chappell (Birmingham, Biosciences), is that evolution "discovered" the power of what we call "meta-configured" genomes -- relatively late products of biological evolution. in a universe with very deep and rich mathematical generative powers, expressed at different stages during individual development, as explained below.

One (unobvious) consequence is that besides including deep mathematical structures, the physical universe includes mechanisms that can be combined to form organisms capable of making increasingly complex mathematical discoveries and and creating or trying to create instances of new mathematical structures.

Such evolutionary mechanisms also require deep mathematical features that have proved hard to understand and to replicate in Artificial Intelligence -- perhaps because they make essential use of sub-neural, chemical, information processes that make essential use of a mixture of discrete and continuous processes that cannot all be replicated on digital computers (including Turing machines).
[This is very different from: "Ontogeny recapitulates phylogeny" (Haeckel)]

DRAFT HIGH LEVEL OVERVIEW

THEMES **Added 6 Sep 2018**

A number of distinct, but intricately related, themes can be found in this project, although this feature was not planned when the project started.
(The order of this list is not significant, and it is likely to be extended/modified.)

- Varieties of forms of reproduction and development, using different epigenetic mechanisms
- Varieties of information processing functions
- Varieties of information processing mechanisms and architectures
- Varieties of types of construction kit
- Varieties of environments and environmental challenges
- Varieties of physical forms and physical behaviours
- Varieties of increasingly complex mathematical structures instantiated in physical

mechanisms, information processing mechanisms, and forms of information (added 4 Mar 2019),

- Varieties of types of compositionality (added 28 Sep 2018):
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html

---

OPEN QUESTIONS ABOUT EVOLUTION

This section updated: Feb-Apr 2018 (Previously 2 Mar 2017)

Closely related questions about biological evolution, listed below, need to be distinguished.

The standard Darwin-Wallace theory of evolution by natural selection, as normally understood, addresses only Question 2(b), attempting to explain how evolution "chooses" between available options. This implicitly assumes that there is an answer to Question 2(a), which asks what makes those options possible. Questions about what sorts of things are possible and what makes them possible are the deepest questions in science, as explained below.

A MAIN QUESTION AND SUBSIDIARY QUESTIONS:
Q.1 What makes various evolutionary or developmental trajectories possible?
There is a huge branching space of possible evolutionary trajectories, only a subset of which are realised at any region of space-time: only a subset of physically possible trajectories become actual trajectories. Each subset spawns new branching possible trajectories --- different subsets in different parts of the universe at different times --- and only a subset of those will actually be realised. A deep theory of biological evolution should explain all the possible trajectories, not merely those that are actualised.

Products of evolution in different portions of space time may differ in physical forms, in behaviours, in kinds of information processing (differences in information available, information acquired, information created, modified and used, information-processing mechanisms used, results of information use, etc.)
What features of the universe make that whole branching space possible?

Q.2(a) What makes new branches possible at different space/time regions of the universe?
Q.2(b) What causes various subsets of those possibilities to be realised in actual evolution (e.g. at different times and different places, or in different individuals, or in the same individual at different stages of development)?

The question "What makes all those evolutionary trajectories possible?", is generally thought to have been answered by the Darwin/Wallace theory of evolution by natural selection. But that theory leaves gaps to be filled by answers to the questions posed here, including questions about varieties of evolutionary possibility that have never been instantiated, but might have been in suitable circumstances. Compare the ability of deep physical theories to predict the consequences of physical configurations that have never been realised. (Much of the creativity in human engineering depends on the existence of previously unnoticed possibilities.)

Q.3 Did some of the trajectories introduce types of information processing that scientists have not yet noticed, that are crucial for understanding natural intelligence? (E.g. ancient mathematical discovery processes?)

I suspect the answer to Q.3 will help us understand why developments in AI so far fail to model or emulate important aspects of human and animal intelligence, including understanding how great mathematical minds made major discoveries long before the development of modern mathematics based on "formal" mechanisms e.g. logic, algebra, set theory, etc. Many of the ancient mathematical discoveries revealed geometrical impossibilities and necessary connections between possibilities, whereas current neural theories focus on probabilities.

Current theories of brain function do not explain abilities to represent necessity and impossibility, although the brains of ancient mathematicians (e.g. Archimedes, Euclid, and many others) clearly used such abilities (as do brains of squirrels and other intelligent non-humans, though they lack the meta-cognitive mechanisms required to reflect on their competences and discuss them). Requirements for such a type of mechanism are being explored here: A Super-Turing Membrane Machine for Geometers (but there's much more detailed work to be done).

Q.4 What features of the fundamental construction kit made all the rest, including derived construction kits, possible? How? Are important kinds of construction kit possible that have not yet come into existence?

Q.5 How are derived construction kits produced by evolution, and how do they change subsequent evolution?

Q.6(a) What sorts of mathematics are required to describe and explain the operation of these construction kits, and their powers and limitations?
Q.6(b) What sorts of information processing mechanism (discrete? continuously variable? deterministic? stochastic?) allowed brains to acquire and manipulate spatial (topological/geometric) information in ways that eventually led to the discoveries of ancient mathematicians, including axioms, constructions and theorems in Euclid's *Elements*?

Q.7 In what ways was evolution a "blind mathematician" using increasingly complex mathematical features of the universe and deriving new mathematical features?

From this standpoint the creativity of evolution exceeds all other known forms of creativity.

(Q.4-Q.7 made explicit here: 27 Dec 2017, Q.2 revised 5 Feb 2018.
Further revisions 28 Feb -- 15 Apr, 2018)

Notes on the questions
Different answers to the first question limit in different ways what can be explained by answers to the second. The third question considers extensions to the answers to both questions, which I suspect Turing might have worked on if he had lived longer. The first and third questions are the main topic of the Meta-Morphogenesis project. The second question is the one most theories of biological evolution since Darwin focus on.

An assumption that pervades all of this is that a scientific understanding of the universe will need to treat at least matter, energy and information as fundamental aspects or constituents of the universe, none of which can be studied independently of the others.

The three concepts cannot be given explicit definitions: they can be explicated only in ever expanding theories about their mutual relationships (e.g. information can control how available energy is used in the behaviours of material things; matter is required for storage of energy and information, all changes in states of matter and information require energy).

Note added 31 May 2018:Turing on mathematical intuition
It was not until 2018, thanks to Francesco Beccuti, that I became aware of Turing's remarks about intuition *vs* ingenuity written in 1938, published in (Turing 1938, Sec. 11). I now wonder whether this distinction was part of the hidden motivation for his paper on morphogenesis:

> *"Mathematical reasoning may be regarded rather schematically as the exercise of a combination of two faculties, which we may call intuition and ingenuity. The activity of the intuition consists in making spontaneous judgments which are not the result of conscious trains of reasoning. These judgments are often but by no means invariably correct. . . . The exercise of ingenuity in mathematics consists in aiding the intuition through suitable arrangements of propositions, and perhaps geometrical figures or drawings."*

Perhaps he was moving toward ideas something like the ideas presented here, including the notion of a super-turing computer with geometric/topological reasoning abilities under development here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-geom.html
But that assumes a connection between his thinking in the late 1930s and his thinking around 1950 for which I have no direct evidence.

The importance of information processing
(Added 5 Feb 2017)
Einstein showed how matter and energy are ultimately the same thing. Information seems to be something different. Information can be used to take control decisions. It can be used to specify goals to be achieved, dangers to be avoided, methods and resources available, preferences between options, requirements for new features of a machine, designs for things that meet the requirements, and many more.

Unlike mass-energy, information does not seem to be conserved, as shown by the ever-increasing variety and power of new kinds of information produced by biological evolution and its products -- using increasingly complex construction kits produced by evolution, combining matter, energy and information.

As the amount of information acquired increases, physical matter in living organisms becomes more organised, including matter used to specify the organisation. So entropy decreases in the relevant part of the universe. However that need not conflict with the law of increasing entropy if there are external sources of energy, e.g. including solar radiation and various forms of energy stored in the planet producing earthquakes, volcanoes, tides and other less dramatic, but biologically important effects. Entropy decrease in some parts is compensated for my entropy increase in others. If a mostly flat tray with many marbles has a collection of deep but narrow depressions (i.e. tubes sealed at the bottom), then if the tray is held horizontal and shaken randomly in a horizontal plane more and more marbles will fall into the tubes, decreasing the randomness of the configuration. Getting the marbles out onto the tray will require carefully deployed energy.

An increasingly large and important subset of evolved construction kits are concerned with information-processing: i.e. acquisition, transformation, storage, combination, analysis, transmission, and above all **use** of information. The earliest uses must have been for simple control functions in the most primitive organisms, e.g. binary decisions, such as whether to move or not, or whether to allow a molecule to pass through a membrane, or homeostatic control, e.g. maintaining temperature. There's more detail on evolved construction kits in Sloman(C-K) especially the crucial role of construction kits for building new kinds of information processing mechanism.

Over time, and across many evolutionary steps and generations, information is used for increasingly varied functions, now including asking questions about evolution. This is reflected in a new alternative title for this project: **The self-informing universe**.

> **Perhaps future science will extend these ideas to fuse the concepts of matter, energy and information in something like the way Einstein's theory of General Relativity fused matter and energy. There's more on this triad below. We need all three to understand what life is and how it evolves -- in a way that produces increasingly sophisticated information-processing machines: minds of many kinds! However, as I keep saying, this concept of information is NOT Shannon-information, but Jane Austen's concept of information. Shannon was concerned with mathematical properties of information bearers, which he unfortunately labelled "information", thereby confusing a whole generation of researchers, though Shannon was not confused, as he was aware of the difference,** explained in the discussion of Jane Austen's concept of information, below.

This introduction focuses mainly on how Q.1 has been ignored by many who support the theory of evolution by natural selection. Later sections will mention Evo-Devo problems relevant to Q2 (a and b) and Q.3.

Originally, the focus of this project was mainly on Q.2 and Q.3 -- the search for important unnoticed transitions in information processing in evolution of biological organisms, transitions whose discovery could give us clues regarding unexplained powers of humans and other animals, e.g. the ability to make mathematical discoveries of the sorts reported in Euclid's *Elements* over 2 thousand years ago.

The importance of Q.1 and the roles of fundamental and derived *construction kits* in the answers to Q.1 did not become obvious until late 2014, although they were implicit in the short paper introducing the M-M project in 2011, published as Sloman(2013). Fundamental and derived construction kits are discussed in more detail in a separate paper:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

Stuart Wray's sketch of aspects of meta-morphogenesis

Anyone who has difficulty taking in all these ideas may find the following helpful. On 5th Jun 2012, Stuart Wray, read a draft conference paper on Meta-morphogenesis and the Creativity of Evolution (before the ideas about construction kits had been added):
http://www.cs.bham.ac.uk/research/projects/cogaff//12.html#1203
To help him understand the contents he produced a sketch of the ideas in the project, reproduced

here, with his permission:

(CLICK THE PICTURE TO SEE A LARGER VERSION)
Wray-MM
CLICK HERE FOR PDF VERSION.
If anyone produces a revised or rival sketch I would be grateful for a copy!

A standard, but incomplete answer

Since Darwin, the standard answer given by scientists (and many philosophers) to Q.1 and Q.2 (above) is "natural selection": instances of alternative designs compete, and some survive and produce offspring, while others do not -- a process that repeatedly produces more complex and more diverse organisms via branching (and sometimes merging) evolutionary trajectories.

But the deeper question underlying Q.1, is often ignored. We can express it thus:

> *What makes it possible for the increasingly complex products of increasingly complex and varied physical mechanisms to come into existence and function in an enormous variety of physical life forms, with new forms regularly emerging?*

Quote from Graham Bell
This question is ignored by many evolutionary biologists. For example, the biologist, Graham Bell, apparently oblivious to Q.1, wrote, in (2008),

> "Living complexity cannot be explained except through selection and does not require any other category of explanation whatsoever."

ignoring the need to explain what makes both the successful and the unsuccessful evolutionary transitions possible candidates for selection. A successful explanation should also account for possibilities that have not been and never will be realised, because the right conditions do not exist, although they are physically possible.

Since the discovery of mechanisms encoding genetic information in DNA, I have the impression that at least some of the researchers who take Q.1 seriously assume that a sufficient answer is provided by postulating mechanisms that produce small random variations in genetic codes. But the possibility of random changes in DNA does not suffice to explain how a planet whose only life forms were single celled organisms could much later produce elephants, giant redwood trees, magpies, and mathematicians like Euclid. Compare someone from another country or another planet asking how pressing buttons on a hot drinks dispenser makes different drinks come out of a spout. Answering that there are selection buttons that specify the desired drink leaves out most of the explanation.

An answer to the question about how evolution produces so many new species over a long time period should also explain possibilities that were not realised (selected) but might have been realised in different circumstances. Just as a deep theory of planetary motion (e.g. Newton's) must explain far more possible trajectories than observed trajectories, a deep theory of biological evolution should explain far more possible products of evolution than observed products. Chomsky 1965 made similar points about the need for "generative" explanations of possible utterances in a language.

NOTE added 26 Jan 2017
I am grateful to John Doyle and Anthony Durity for drawing my attention to the work of Kirschner who clearly recognized this problem, confronted in M.W. Kirschner and J.C. Gerhart (2005). See also this interview with Kirschner
http://bmcbiol.biomedcentral.com/articles/10.1186/1741-7007-11-110
"So I think that to explain these developments in terms of the properties of cell and developmental systems will unify biology into a set of common principles that can be applied to different systems rather than a number of special cases that have to be learned somehow by rote."

This seems to be consistent with the theory of evolved construction kits being developed as part of the M-M project.

AISB 2017 paper
A progress report on part of the M-M project was presented at the Symposium on Computing and Philosophy at the AISB 2017 Convention, Bath University, April 2017
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-aisb17-CandP.pdf

Related work in Biosemiotics

Added 16 Jan 2018:
Alexei Sharov drew my attention to deep, closely related work by the **Biosemiotics** research community, e.g.:
http://www.biosemiotics.org/biosemiotics-introduction/
https://en.wikipedia.org/wiki/Biosemiotics https://en.wikipedia.org/wiki/Zoosemiotics
I shall later try to write something about the connections, and will add links on this site.

JUMP TO CONTENTS

Natural Selection’s Explanatory Gap

The only thing natural selection explains is why some of the temporarily realised possibilities do not survive. They don’t survive because rival organisms (or rival designs) cope better in the environment at the time, and for some reason alternative designs cannot be supported, e.g. because of scarce resources. (There’s much more to be said about this, but not here.)

But such "standard" answers do not explain what makes those rival life forms possible: it is just assumed that somehow they come into existence and when there is competition for resources some forms are better able to obtain resources, thereby depriving the alternative forms, so that they die away. But that doesn’t explain what is it about our universe that makes all those enormously varied life forms possible -- along with others that might have evolved but did not, since the processes are random?

This question was taken seriously by Erwin Schrödinger: (1944), who provided a partial answer, as explained below.

At this stage some readers may be reminded of theoretical biologists, e.g. Brian Goodwin, who have emphasised "laws of form" rather than genes. I am not aware of any version of that theory that has sufficient explanatory power, but it can be considered a step in the right direction.

Explaining possibilities

Deep explanatory theories must have deep explanatory powers. The important role in science of explanations of possibilities (as opposed to regularities/laws) was discussed in Chapter 2 of Sloman (1978) (expanded slightly in the latest edition), but is not yet widely understood. For example, it contradicts Popper's claim (modified in his later work) that good scientific theories must be falsifiable empirically.

A loose analogy may help to clarify what I am saying. If people are observed playing a game, someone watching might want to know (1) what moves are possible according to the rules of the game and the capabilities of the players, and which impossible, either because of the rules of the game or limitations of the players or aspects of the environment.

The observers may also wish to know (2) how the players choose among the possibilities. The answers to (1) could refer to different sorts of enabling and constraining factors, including rules of the game and capabilities of the players. The answers to (2) could refer to mechanisms that operate within the possibilities and constraints, including both mechanisms for selecting between possibilities and mechanisms for realising and deploying the selections.

Since Darwin, evolutionary biology has made great strides in elaborating answers both to Q.2 and to some aspects of Q.1, although most methodological discussion of evolutionary explanations tends to focus on which among competing lineages wins the battle for resources, often without explaining what mechanisms make the competing lineages possible, leaving Q.1 unanswered. Q.3 is usually not even considered, in my experience.

Most of the progress on Q.1 is concerned with mechanisms related to physical forms and physiological processes. Evolution of mechanisms required for *information processing capabilities* (e.g. perception, learning, reasoning, explaining, development of new concepts, etc.) has received much less attention. This may in part be explained by the enormous difficulty of the task (partly because information processing mechanisms do not generally leave fossil records) and in part by the fact that most scientists (including most psychologists and neuroscientists) receive a very shallow education regarding information processing mechanisms -- most (but not all) merely learn to use a small subset of tools designed by others, especially tools for analysing and evaluating research data rather than tools for designing and testing explanatory mechanisms, e.g. visual mechanisms, mathematical reasoning mechanisms, and mechanisms required for use of human languages.

So the researchers do not acquire deep competences in what was referred to as "Computational Thinking" in Wing (2006) (echoing older ideas of Alan Kay, Seymour Papert and Marvin Minsky). Unfortunately, current dramatic changes in computing education do not seem to be addressing this gap, since they focus mainly on *practical* uses of computers, or uses for entertainment, not deep *scientific* explanatory uses.

The current state of the art in the science and engineering of information processing systems is still not up to the task. As far as I know, no existing scientist (or engineer) could design a working brain for a robot that interacts with its environment as a human toddler, or a monkey, squirrel or crow does and also develops new information processing competences as they do. Answering that aspect of Q.1 remains a deep challenge.

An illustrative explanatory gap

Soon after I started learning AI programming around 1970 I set myself the goal of working out how to design a "baby robot" that could grow up to be a mathematician able to make discoveries in Euclidean geometry, hoping to test some of Immanuel Kant's ideas about mathematical knowledge in Kant(1781), whose views I had defended in my 1962 DPhil thesis, before I knew anything about AI.

Nearly half a century after proposing that goal, I still know of no AI researcher who has achieved the goal or is close to achieving it (wittingly or unwittingly), and the reasons for lack of progress are far from obvious, though one of the reasons is failure to notice the problem.

My main motivation for the Meta-Morphogenesis project is the hope of discovering important clues about what is missing in current theories of cognition -- especially theories about spatial cognitive abilities used by ancient mathematicians long before Euclid. For various reasons, including ignorance of ancient geometric discoveries brought about by misguided changes in mathematical education policies during the 20th century, very few children now encounter geometric proofs and constructions that used to be a standard part of mathematical education. As a result, hardly any psychologists, neuroscientists, cognitive scientists, and AI researchers think about cognitive mechanisms required for making those ancient mathematical discoveries.

A major aim of M-M project is to fill that gap by studying previously unnoticed aspects of evolved biological information processing systems, in the hope of answering Q.1 (What makes various evolutionary or developmental trajectories possible?) Many researchers just assume, like Graham Bell, quoted above, that new designs are somehow produced by biological mechanisms and inferior ones eliminated by competition for resources.

In contrast, a key idea in the M-M Project (partly inspired by Turing's work on morphogenesis [Turing-52]) is that Q.1 needs an answer, and a key part of the answer, developed in a separate paper Sloman(C-K), is that the physical universe supports a wide variety of construction kits derived from a "fundamental" construction kit (FCK) provided by physics and chemistry, or possibly some deeper but still unknown physical mechanisms underlying physics and chemistry.

The hope is that the theory of fundamental and derived construction kits will unify physics, chemistry, biology, and eventually also AI, neuroscience, psychology and philosophy, providing new answers to many old philosophical questions about the nature of life and mind (including the powers of mathematical minds).

While Q.1 remains unanswered, evolution of animal intelligence, including human intelligence, remains unexplained: the Darwinian answer in terms of reproductive fitness does not explain what makes the observed kinds of animal intelligence possible. It merely explains how IF intelligent species evolve THEN they are more likely than less intelligent competitors to survive in environments where intelligence is useful. Of course less intelligent species with which they are not in competition may also survive, e.g. microbes, insects, plants of many kinds, etc.

I suggest (partly inspired by Turing's work on morphogenesis [Turing-52]) that Q.1 needs an answer, and a key part of the answer, sketched but with many remaining gaps, in a separate paper Sloman(C-K), is that the physical universe supports a wide variety of "construction kits" derived from a "fundamental" construction kit (FCK) provided by physics and chemistry, or possibly some deeper but still unknown physical mechanisms underlying physics and chemistry. Some of the construction kits are concrete, others abstract, especially construction kits for building information processing systems, as illustrated by the ever-expanding uses of virtual machinery in computer based systems over the last half century -- a transition whose biological relevance has mostly gone unnoticed. There are many things the fundamental construction kit cannot produce without first producing many layers of derived construction kits and their products, including microbes, monkeys, mathematicians and musical boxes.

The M-M project was launched (in 2012, unfunded) in the hope that by investigating previously unstudied evolutionary transitions we may gain clues regarding forms and mechanisms of information processing that we have not stumbled across in psychology, neuroscience, computer science, software engineering or Artificial Intelligence, including some that will provide clues as to mechanisms required for replication of mathematical competences that made possible the discoveries organised in Euclid's Elements, and related competences in human toddlers and other animals that we cannot yet explain or replicate.

The label for the project "Meta-Morphogenesis" was based on the observation (which could not have escaped Alan Turing) that many of the products of evolution (a form of morphogenesis) alter the mechanisms of evolution, partly by altering the opportunities and tasks faced by products of evolution.

A simple example is the evolution of sexual reproduction, which was not possible in the earliest (unicellular) life forms. A more subtle example is the evolution of mate selection making use of previously evolved cognitive mechanisms that allow possible mates to be compared and evaluated, modifying the selection processes and thereby the course of evolution (sometimes in bizarre ways). There are also many ways in which evolution of one species can affect evolution of another, e.g. when one species uses another for dispersal of its seeds.

More generally, evolution of certain cognitive powers enabled formation of cultures and evolution of cultures, which in turn could influence natural selection. A possible result that has not, as far as I know, occurred might be evolution of tendencies towards religious belief that in some societies might lead to oppression of those who lack the tendencies, including preventing them from producing offspring. This might lead to a culture dominated by individuals genetically disposed to superstitious beliefs (if such a thing is biologically possible).

The importance of derived construction kits

Products of evolution include not only new species, and new variants of old species, but also new construction-kits: "derived" construction kits (DCKs). Some of them are concerned only with uses of physics and chemistry to create new physical/chemical structures (e.g. muscles, skeletons and exoskeletons), whereas others are used to create new information-processing mechanisms, serving increasingly complex functions as evolved organisms became more intelligent.

Each construction kit provides materials and mechanisms that support a particular branching space of possible sequences of assembly steps, where each step both consumes resources, thereby limiting future possibilities and also supports new possibilities by combining old mechanisms to create new structures and mechanisms that can also be combined.

In some cases temporary structures (e.g. scaffolds or tools) are produced that allow new assembly processes to occur, after which the temporary structures are discarded or disassembled. Disassembly can open up further new branches based on re-use of discarded parts.

Of particular importance for biology are the increasingly complex new construction kits for building increasingly sophisticated information-processing systems (eventually including human minds).

Readers who have had experience "playing" with construction-kits such as Meccano, Tinker Toys, Fischertechnik, plasticine, sand, paper and scissors, etc. may have the advantage of finding it easy to think about a universe providing a fundamental construction kit from which everything else in the universe is built. These ideas can be elaborated in ways that many will find surprising.

Unlike familiar "toy" construction kits, the fundamental construction kit provided by the physical universe somehow continually produces increasingly complex and varied construction kits that grow out of previously evolved construction kits (as discussed in more detail in a separate paper Sloman(C-K)).

Unlike many physical/chemical processes in which patterns form, interact, and change, evolutionary processes make essential use of **information**, e.g. genetic information specifying and controlling the construction of new offspring and acquired perceptual information used in selecting and carrying out increasingly complex actions.

As evolution produced organisms with more varied and complex developmental and behavioural competences, they needed to **use information** to choose between competences to use at various stages of development and when acting in complex environments.

The types of information, uses of information, forms of information representation, and mechanisms of information processing became increasingly complex and varied over time. The crucial fact about information is that it is used in control processes, including control of reproduction, development, interaction with the environment, mating, social interaction, and collaborative constructions. Since reproduction and evolution make essential use of information, this meant that the mechanisms of reproduction and evolution could be changed by evolution. Hence the label "meta-morphogenesis".

Unfortunately, the study of evolved forms of information and mechanisms of information-processing can be much harder than the study of evolved physical forms and physical behaviours, which leave fossil records, tools, and other artefacts. Was Turing thinking along these lines when he died?

CONTENTS

Hints from Alan Turing

As Turing seems to have realised: the forms of information-processing used by evolution were richer and more varied than those developed by computer scientists and engineers so far, and made essential use of chemistry.

> **"In the nervous system chemical phenomena are at least as important as electrical"**
> Turing (1950)
>
> That comment is almost universally ignored, though it occurred in one of his most widely cited papers.
>
> Re-reading it in 2011 led me to wonder, especially in the context of trying to write some comments on his 1952 paper, whether he had thought about the significance of chemistry for evolution of information processing mechanisms rich enough to support minds in a physical universe.

Two years after that he was dead. What would he have done if he had lived several more decades?

Is this a clue to the answer:

> In 1946 Turing wrote a letter urging W. Ross Ashby to use Turing's ACE computer to implement his ideas about modelling brains. He expressed a view that is unfashionable among AI researchers at present (2016), but accords with the aims of this project.
> He wrote:
>
> > *"In working on the ACE I am more interested in the possibility of producing models of the actions of the brain than in the practical applications to computing."*
> > http://www.rossashby.info/letters/turing.html

Perhaps Turing would have worked on the Meta-Morphogenesis project, namely

> Trying to understand the many steps in evolution of increasingly complex biological information-processing systems: including molecules, microbes, mastodons, magpies and mathematicians.
>
> This could lead to discovery of previously unnoticed mechanisms in biological brains and minds that explain the many discrepancies between animal intelligence and current AI, and perhaps also explain how new forms of intelligence not yet produced by natural selection are possible even though biological evolution has not yet stumbled across them.

NOTE added 12 Mar 2015:
A different answer to the question about Turing, by Andrew Hodges.
It turns out that in 2002 Andrew Hodges asked exactly the same question about what Turing might have done if he had lived longer. He gave a completely different answer, without mentioning Turing's interest in biology. His answer was published here in 2004:

What would Alan Turing have done after 1954?
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/m-m-related.html#hodges-turing

BACKGROUND: How did this project start? In Memoriam: S. Barry Cooper (1943-2015)

Added 10 Nov 2015; Updated 15 Jan 2017; 11 Mar 2018
This project resulted from a miscommunication in 2011 with Barry Cooper, who had asked me to contribute to a book he was co-editing, on Alan Turing's life and work, eventually published in 2013, the year after Turing's centenary: Cooper and van Leeuwen, (2013) (winning several awards). Further information about the book and a full table of contents can be found here: http://www.cs.bham.ac.uk/~axs/amtbook/.

This document started as a description of my contribution to that book, especially my paper in Part IV, proposing the Meta-Morphogenesis project. That project later spawned a paper on fundamental and evolved construction kits, mentioned below.

Alas, I learnt in November 2015, that Barry had died unexpectedly after a very short illness. At that time he and Mariya Soskova were in the process of editing a Book *The Incomputable*, to be published by Springer, to which he had invited me to contribute another paper extending the 2011 paper defining the M-M project. The new paper is on the crucial roles of various sorts of construction kit in biological evolution. A version of that paper frozen in 2016, was published in 2017 as an invited contribution to Cooper and Soskova(2017). A revised, extended, still growing, version of that paper is also available Sloman(C-K).

For further information about Barry Cooper and links to early tributes see
http://www.math.uni-hamburg.de/spag/ml/ACiE/

My personal debt to Barry Cooper
Barry and I had never met until we were both invited to contribute chapters to a book published in 2011 on *Information and Computation,* edited by Gordana Dodig-Crnkovic and Mark Burgin, and published (at an exorbitant cost, by World Scientific):
http://www.worldscientific.com/worldscibooks/10.1142/7637

Barry and I reviewed each others' chapters and as a result had some direct communication for the first time, by email -- producing mutual respect. Later, in 2011, out of the blue, he invited me to contribute to a proposed Turing centenary volume originally described here:
http://www.mathcomp.leeds.ac.uk/turing2012/give-page.php?300

After some discussion about how best I could contribute, which involved a change of plan mid-way, I sent him my promised three related short chapters, the first for Part I, of the book (How Do We Compute? What Can We Prove?) and the others for Part III (Building a Brain: Intelligent Machines, Practice and Theory).

However, as a result of the earlier change of plan, which Barry had apparently forgotten, I was still listed as a contributor to Part IV (The Mathematics of Emergence: The Mysteries of Morphogenesis). So he later asked me for my missing contribution to Part IV. His request led me to look at Turing's paper on morphogenesis published two years before he died, Turing 1952

That triggered questions about what Turing might have done if he had lived another 30-40 years, instead of only two. My speculative answer to that question identified the Meta-Morphogenesis (M-M) project, so-named because biological evolution is a process of morphogenesis that modifies the mechanisms of morphogenesis. The key focus of the M-M project is on changes in information processing mechanisms produced by biological evolution. I conjecture that there remain important mechanisms that have not been discovered and which are essential for understanding and replicating important brain functions in humans and other animals, especially those concerned with making mathematical discoveries.

So I offered Barry a short paper proposing "The Meta-Morphogenesis Project" as a conjectured answer to my hypothetical question about Turing, written in 2011, published as Sloman(2013). He accepted it (as the final commentary paper in the book) and ever since then I have been working full-time on the project, for which this page is the main entry point, from which many sub-projects have been launched.

Barry later invited me to talk to a workshop at the Royal Society Kavli Centre in June 2012, followed by an invitation to talk to the Logic Seminar in Leeds in October, and then in 2013 invited me to contribute to a book *The Incomputable*, to be published by Springer, being co-edited by Barry and Mariya Soskova, mentioned above.

My chapter, finished in 2016, not long after Barry's death, extends the idea of the Meta-Morphogenesis project by discussing requirements for many kinds of construction kit to be used in biological evolution, including construction kits for producing
-- physical components of organisms,
-- information processing systems,
-- new construction kits.

Some of the construction kits are **concrete** e.g. mechanisms for building components of tree trunks, while others are **abstract**, e.g. construction kits for reasoning mechanisms implemented as virtual machines.

A "work in progress" version extending the published work on construction kits is online here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
An earlier version frozen mid 2016. was included in the new volume, published in 2017.

I also met Barry a few times at workshops or conferences, and I always immensely enjoyed and benefited from our discussions. Originally stimulated by Barry's requests and enormously encouraged by his responses to my drafts, I now expect to go on working on these problems as long as I can.

So he changed my life, by giving me a new research direction, which does not often happen to 75-year old retired academics! (Now even older.)

I am very sad that we'll not be able to have any more conversations.

Formulating the goals of the Meta-Morphogenesis project made me realise that much of my own work of the last 40 years could be re-cast as a contribution to that project. This document, and a growing set of linked documents is my (messy, changing, inept) attempt to present the project: its questions, some of what we don't know, some of what we do know, some of the ways we can make progress, and some of the overlaps with work of other thinkers. (I suspect Immanuel Kant was attempting to work on these topics in his *Critique of Pure Reason* (1781), but lacked important conceptual tools developed in the 20th century, connected with information processing systems.)

> **I am well aware of the risk, earlier pointed out by Erwin Schrödinger:** (1944)
> (a book that can be seen as an important contribution to the M-M project):
> *"I see no other escape from this dilemma (lest our true aim be lost forever) than that some of us should venture to embark on a synthesis of facts and theories, albeit with second hand and incomplete knowledge of some of them - and at the risk of making fools of ourselves.*
> *So much for my apology."*
>
> Me too! (A.S)

I do not intend to apply for any funds for this project. Others may, if they wish.

CONTENTS

Meta-morphogenesis: a more detailed overview.

The possibility of life depends on a very powerful "fundamental" construction kit provided by the physical universe, i.e. physics+chemistry, including features not yet understood, and using natural selection (plus serendipity) to produce:

- many branching layers of increasingly complex physical/chemical structure and mechanism -- different types in different parts of the universe.
- many branching layers of *information-processing* machinery, required for new forms of control, new forms of reproduction, new forms of development, new forms of intelligence, new forms of social/cultural evolution, continually producing and using new types of concrete and abstract "derived" construction kit.

Information is primarily something to be used not transmitted

It is widely assumed that information is essentially something to be *transmitted* from a sender to a receiver, though for various reasons it may also need to be *stored*, *transformed*, or *combined* with other information.

But being stored and being transmitted cannot be the most important processes and states involving information, since information is acquired, stored, transmitted and received *because it can be used*.

So an adequate theory of information, including biological information, must include an account of the *uses* of information, not just the syntactic forms of information-bearers or mechanisms of coding, decoding, error correction, etc. (I think Shannon understood this, unlike many of his admirers.)

Paragraph revised 27 Jul 2018
I previously thought that Gregory Bateson’s much quoted but bizarre claim that "information is a difference that makes a difference" is uselessly vague -- perhaps an attempt to draw attention to the fact that the basic type of use of information is for *control* (making a difference to what

happens), and everything else flows from that.

> What Bateson might have meant is discussed further in:
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/information-difference.html

Control takes many forms, with many different information contents, and it uses many different sorts of information. New contents, forms, uses, and mechanisms are constantly being found/created by biological processes including development, learning and evolution, including cultural evolution.

Storage and transmission are important only because of gaps between when and where information is available and when and where it can be used.

The tasks of understanding uses of information, forms of information, storage of information, transmission of information, transformations of information are major themes of this project.

The novelist Jane Austen seems to have understood that information is fundamentally something to be *used*, a century before Shannon's work, as demonstrated below.

The study of the uses of information in organisms and machines was of great interest to Turing, and also a major theme of Cyberneticians such as Norbert Wiener, W. Ross Ashby, and W.T. Powers ("Perceptual control theory"). But most such researchers tended to assume that all the mathematical structures and processes could be described in essentially numerical terms, using equations (e.g. including differential equations). That is a disastrous limitation, as shown by the work of linguists, software engineers, architects, mathematicians (especially the pioneering ancient mathematicians whose work was assembled by Euclid), and many others who study mechanisms involving creation and use of non-numerical structures.

The "General systems theory" of Bertalanffy may have been an exception to the assumption of numerical information processing, because of his emphasis on structural change and growth: I don't yet know enough about the theory to be sure. In particular, I suspect his theory was not rich enough to accommodate growth of multi-layered networks of *virtual* machinery, an idea that came to be understood by engineers and some others, in the last half century, though so far merely scratching the surface -- a topic discussed in connection with Virtual Machine Functionalism (VMF) below.

The crucial roles of mathematics in evolution

Natural selection repeatedly implicitly discovered and used new types of mathematical abstraction, especially abstractions with non-numerical parameters that could vary in many ways -- long before there were any human mathematicians able to think about mathematical abstractions. Examples included structural parameters (e.g. a grammar, or a biological function, could be a parameter). See also the Multiple Foundations paper.

All of this depended on the ability of the Fundamental Construction Kit (FCK) to be capable of producing a (never ending?) branching and merging collection of new, increasingly complex and powerful, derived construction kits (DCKs), with new mathematical properties, some of which are adopted, or created, by new more complex organisms with more complex functionalities, and needs.

Not all the derived construction kits used by natural selection are products of natural selection: some are products of non-biological physical, geological, astronomical or cosmic events and processes. Evolution can use effects of volcanoes, tides, seasons, asteroid impacts, and cosmic radiation.

Some limiting cases of construction kits may be better thought of as scaffolding: mechanisms and structures that are used *temporarily* in production of new mechanisms and structures -- like diggers, cranes, computers, information stores (including plans) and workers -- used during construction of a sky-scraper, but not in the finished product.

Compare the chrysalis/pupal stages in some insect forms mentioned above, and the role of the placenta in mammals.

CONTENTS

Different (partly overlapping) themes/sub-projects (Skip this list on first reading) (DRAFT Added 28 Feb 2018)

Since 2011, several different (partly overlapping) themes/sub-projects have emerged. This is a messy, incomplete, inaccurate attempt to "pull out" some of the higher level themes added between 2011 and 2018. with many overlaps, at different levels of abstraction. It will need substantial editing as my thoughts clear and I find time.
The topics include evolution of:

-- new physical life forms (physical size/shape/morphology/behaviours), using
-- new, more complex, types of physical/chemical structures and mechanisms,
-- new mechanisms producing and controlling growth and development
-- types of use of information in reproduction and development,
-- types of internal and external information acquired, stored, used, shared, etc.
-- new mechanisms for control of internal/external behaviours
-- types of internal information-processing functions and mechanisms
-- uses of information in controlling internal physical/chemical processes
-- changing requirements for and uses of meta, meta, ... information,
including information about uses of different forms of information,
e.g. topological, geometrical, numerical, grammatical,...
-- uses of information to generate, compare, schedule, modify,
motivators, preferences, goals, plans, etc.
-- uses of information in controlling external physical/individual/group behaviours
-- changing requirements for internal information processing
-- changing requirements for information about the environment,
including other information users,
-- new types of motivation, and mechanisms for motive-generation, motive comparison,
motive selection, motive postponement, meta-management, etc.
-- new opportunities or requirements for re-scheduling, re-evaluating, or abandoning
some current motives, plans, or schedules.
-- new types of self-directed and other-directed meta-cognitive/reflective control,

and adding new layers of meta-cognition (varieties of consciousness),
-- new reproductive mechanisms and processes: vegetative, asexual, sexual, parasitic, cultural ...
-- new physical/chemical/biological/... environments/niches, producing -- new requirements for control of internal/external behaviours, and reproductive processes
-- new meta-mechanisms for creating information-processing mechanisms in reproduction, development, learning, ...
-- new forms and mechanisms of representing, encoding, manipulating information, increasingly sophisticated mathematical (geometric, topological, ...) reasoning capabilities, increasing discovery and use of powerful re-usable mathematical abstractions.
-- forms and mechanisms of information representation/encoding and their uses
-- spatial and temporal scope of biological information processing (e.g. history, prediction, planning, reference to remote locations, future possibilities)
-- new types of construction kit for creating physical chemical structures
-- new types of construction kit for creating new construction kits -- (meta-construction kits)
-- new types of information-based cooperation and competition
-- processes and mechanisms of cognitive/information-processing development
-- other-directed meta-cognition (offspring, conspecifics, competitors, prey, predators
-- types of sharing of information and information processing (e.g. teaching, collaborating)
-- external shared forms of information storage and use (Popper's "Third World")
-- uses of information in designing, creating, using, clothing, tools, shelters, ...
-- uses of information in designing, creating, using, physical machines and other artefacts
-- formalisms and concepts able to express questions, theories, predictions, designs, explanations etc.
-- theories about what exists, how things work, what could exist, the rest of the universe
-- types of explanatory theories
-- oppressive uses of information and misinformation (e.g. myths, religions, tyranny)
-- constructive uses of shared information for mutual benefit (cultures, schools, libraries, research institutes...)
-- the Meta-Configured Genome
-- multi-species food-chains, resource-chains, entropy-chains, (chains include cycles)

All the above need to be revised, extended, reorganised, pruned, etc.
More details are in referenced papers developing these themes.

REQUIREMENTS for organisms and their components can also evolve.

A biological *niche* is something like a collection of requirements, not a geographical location. *The space of possible designs* for functioning organisms and *the space of possible niches* have complex relationships. A change in *design* for organisms of type T1 can affect changes in *niches* (sets of requirements) for organisms of type T1, and others T2, T3, etc.

Thus there are complex relationships between trajectories in different parts of design space and niche space (see Sloman (1995) and Sloman (2000)).

Interacting evolutionary and developmental trajectories of different types on different spatial and temporal scales relate to changes in: whole organisms, parts of organisms, behaviours of organisms, competences of organisms, types of information used, whole species, ecosystems, scaffolding, physical/chemical mechanisms, information-processing mechanisms, and also changes in the requirements for various mechanisms.

All of these changes proceed in parallel, with many interactions between them including production of *new sets of requirements* (new niches).

Some of the newly emerging requirements are physical: requirements for new kinds of structure, new kinds of energy store, new forms of locomotion, new kinds of chemical synthesis, new biochemical defence mechanisms, new forms of motion, and many more. Others are new requirements for information processing.

For some related ideas see Steve Burbeck's project on life and computation.

Other changing requirements concern *information*

New kinds of useful information that become available e.g. because of physical changes produced by geological or other physical changes, new ways of acquiring information (e.g. about the environment, or about internal states and processes), or new kinds of information (e.g. information about information), or new ways of storing information or new ways of processing information, and new ways of communicating information.

A consequence of these ideas is that the common view that mathematics is a product of human minds is back to front. All life forms and all important features of life forms, including human minds, are products of the mathematical generative potential of the physical universe interacting with biological evolution by natural selection. New products continually produce new potential. Natural selection would not suffice without that fountain of potential new products. (Recursion was at work long before humans discovered recursion.)

Because many information processing tasks require construction and modification of complex structures with speeds, variety and complexity that could not be achieved by physical mechanisms, the above processes must somehow have produced powerful new virtual machinery long before human engineers discovered the need for such things (in the 20th Century).

Consequences of use of construction kits

One of the consequences of the construction kit idea is that the kinds of mathematics required for fundamental physical theories and production of new physically testable predictions may not suffice for describing features, products, and behaviours arising out of new derived construction kits: this is already evident in the mathematical differences between notations of fundamental physics and notations found useful for describing complex chemical processes in which structures and processes change. A complex molecule like haemoglobin or DNA is not usefully describable merely in terms of a collection of numerical measures.

> There are similar changes in kinds of mathematics required for describing *parts* of meccano sets and those required for describing increasingly complex meccano *constructions*, in which new kinds of constraint on relative motion emerge. E.g. one perforated strip is rigid. If two rigid strips are joined by a single screw (with nut left loose) going through a pair of end holes, a flexible structure emerges, with a changeable angle between the strips. Adding another strip and screw produces yet another kind of structural flexibility since there are two angles that can vary independently. Joining the two remaining free ends with a screw removes all that flexibility: the shape is fixed -- an important discovery in Euclidean geometry. That rigidity

could be an important aspect of the kind of constrained flexibility of a larger structure containing the triangle, as any mechanical engineer will know.

Likewise there are many information structures required for coping with complex environments that instead of the mathematics of arithmetic, calculus, probability and statistics need the mathematics of topology, grammars, semantic contents, logical relationships -- and probably kinds of mathematics not yet discovered. Some examples have already been discovered by engineers whose products have multiple stable and unstable states with different energy relationships and different sorts of constraints on trajectories between states. Natural selection made many such discoveries long before humans did.

Yet more kinds of mathematical complexity emerge from interactions between concurrently active virtual machines, some partially controlling others. (The phenomena included in the study of chaotic dynamical systems are trivial by comparison: they can't produce a mind that discovers a theorem in topology -- as toddlers seem to be able to do, unwittingly.)

It is remarkable that the FCK had multi-layer self-extending capabilities able to produce physical and virtual machines with information processing capabilities that eventually generated questions about the nature of the universe and its ability to produce life.

But many of the intermediate stages and parallel branches are equally remarkable.

The need for a theory of construction kits

Without a theory of the powers of the fundamental construction kit, supplemented with many theories about the derived construction kits that the FCK is able to support, the theory of evolution by natural selection is based on an unfulfilled promise of explanatory power.

IN OTHER WORDS: Without a detailed theory of the FCK and the many sorts of DCKs produced from it by combinations of natural selection and physical/cosmic accidents, the theory of evolution by natural selection is something like a story about the many travels of a vehicle that never mentions the engine of the vehicle.

When we know a lot more about the FCK there will be many surprises -- none of which will require an intelligent designer however.

The processes of evolution by natural selection have been compared with the work of a blind watchmaker. I prefer the comparison with a "blind mathematician" stumbling across theorems about possible products of the fundamental construction kit. The proofs of the theorems of possibility are the evolutionary trajectories leading to instances of those possibilities. But only some of the most sophisticated products of evolution have any hope of understanding the theorems or their proofs. We still have a long way to go.

This document is a product of products of physical and evolved construction kits, including social/cultural construction kits.

Fundamental and derived construction kits

There is more on the roles of fundamental and derived construction kits in evolution in this messy draft, incomplete document, begun in Nov 2014: Sloman(C-K). (Some of the ideas will appear in a Springer book chapter, currently in press.)

All of this raises questions about whether the current forms of mathematics used by physicists have rich enough generative power to explain the generative power of the FCK, a type of question raised in the context of reproductive mechanisms by Schrödinger (1944). Compare the challenge posed by a Physics Nobel laureate Anderson (1972). I suspect our educational system needs to include far more varieties of mathematics, including the recently much neglected mathematics in Euclid's *Elements* (published about 2,500 years ago, and arguably the single most important publication on this planet). The new educational system will also need to include an introduction to the variety of forms of virtual machinery, their mathematical properties, and their causal powers. I don't mean only virtual machines implementable on Turing machines or digital computers.

CONTENTS

Two major themes: Construction kits and mathematics
(Added 5 May 2015 - revised Aug 2015)

In November 2014, while writing notes for talks in Edinburgh and Turin, I realised that a central idea in this project must be the idea of a "construction kit", explored in a sub-project introduced in Sloman(C-K).

Evolution initially uses the "Fundamental" Construction Kit (FCK) provided by Physics/Chemistry, but repeatedly extends the scope of natural selection by producing ever more complex and powerful "Derived" Construction Kits (DCKs). Some of the DCKs are concrete (physical), some are abstract, and some are a mixture: hybrid construction kits. The processes also require scaffolding: use of temporary constructs. Many of the later construction kits are concerned with production of virtual machinery: about which human engineers have learnt a great deal since the 1950s, though the new knowledge has been largely ignored by philosophers (and neuroscientists?), and the philosophical, psychological, and biological implications have been ignored by most computer scientists, biologists, psychologists, and engineers.

> **Note:**
> This is a very unfortunate communication gap (not being addressed by the new wave of enthusiasm for teaching computing to children -- a missed opportunity, which may have to wait for a new generation of computing teachers with broader interests and backgrounds).

Note on Virtual Machine Functionalism

A "high level" tutorial introduction to the some of the ideas about virtual machines is here:

> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html
> Virtual Machine Functionalism (VMF): The only form of functionalism worth taking seriously in Philosophy of Mind and theories of Consciousness.
> Some implications for the study of consciousness are discussed in Sloman and Chrisley 2003.

Construction kits intrinsically have mathematical properties: their generative powers and their constraints/limitations. That's why, to a large extent, it is easy to recognize and distinguish novel objects produced using a Meccano kit, a Lego kit, or a Tinker-toy kit. Because the generative powers, in most cases, are recursive or iterative, in the sense that constructed products of a kit can always be extended using that kit, and all the products of a kit are constrained by properties of the kit, the processes of evolution are deeply mathematical. However, when products of a construction kit CK interact with things that are not solely products of CK, the resulting processes can have features that are not derivable from properties of CK.

The theorems of evolution (the blind mathematician) are primarily theorems about what is possible including construction kits and their products. The proofs of the theorems are implicit in the evolutionary and developmental trails leading to instances of the possibilities.

Evolution also uses theorems about what is impossible: constraints on possibilities, i.e. necessities. But these do not have proofs produced directly by evolution, until some of the late products of evolution begin to notice them, think about them, reason about them, and communicate the results of such processes.

All this is part of a long and complex story that may not become clear for many years.

The need to provide a foundation for all the products of biological evolution may one day turn out to imply previously unnoticed constraints on the Fundamental Construct Kit (FCK), with important, new, implications for theoretical physics. At the very least this may help to resolve disagreements about fundamental physics. Alternatively it may reveal gaps in physical theory that require major changes. I expect new kinds of mathematics will be required, in order to bridge the gaps between fundamental physics and some of the complex products of evolution. (Compare Anderson 1972.)

Limits to reductionism
Although the products of evolution are fully implemented in physical mechanisms, accurate descriptions of the powers and actions of some of the new virtual machines require use of an ontology that is not definable in terms of the ontology of physics (e.g. concepts of winning, losing, attacking, defending, noticing a threat, trying to construct a defence). It follows that some of the properties of the virtual machines, and descriptions of some of their states, cannot be derived logically (or mathematically) from descriptions of the underlying physical machines. That non-derivability is consistent with full implementability. This is the variety of anti-reductionism proposed in Sloman (1978, revised). It is also one of the main themes of the theory of Virtual Machine Functionalism (VMF) mentioned above.

CONTENTS

The centrality of information
(Added 29 Oct 2014)

(Control information, referential information, how-to information, explanatory information, information about information, various kinds of self-directed information, and many more.)

In contrast with the majority of evolutionary research (that I know of), this project focuses on changes in types of information and types of information-processing in evolution. Those changes produce changes in the roles of information in control, reproduction, development, discovery, learning, communication, coordination and other processes, in living things of all sorts. But the examples keep changing, and becoming increasingly complicated, as a result of evolution.

I am not sure whether there is any well-defined upper bound to the complexity, though if there is one it is likely to be far beyond the types of complexity found so far on earth. I shall not discuss the implications of that except to note that there's no reason to believe evolution of information processing has stopped, or is close to stopping, not least because changes in physical designs for organisms, i.e. genomes, are not required for evolutionary changes in information processing: as shown by cultural evolution, including evolution of art and science, and most recently the internet and related technologies.

Is virtually unending change in information processing an inevitable consequence of the existence of the universe? I don't know. It would not happen in empty space. It would not happen on a planet derived only from grains of sand. If a planet, or solar system, or galaxy, or universe has enough diverse chemical components and enough random influences, then perhaps the unending (frequent? infrequent?) initiation of processes of evolution, including various types of meta-morphogenesis, is inevitable -- though I don't know how constrained the set of possible evolutionary trajectories is. Not even natural selection can produce chemically impossible brain mechanisms!

Perhaps, when we have a deeper understanding of the Fundamental Construction Kit provided by physics and chemistry and its role in biology, we'll be better able to discuss whether a different kind of universe might have support life in a comparable variety of forms.

It is not clear whether the variety of forms of information processing currently known to science (and engineering) can support the variety of possibilities required to support all the products of natural selection. The Church-Turing thesis, if true, may turn out to be true only of a class of computations that can be performed on numbers (plus structurally equivalent computations). We'll see that there are many forms of information-processing that are not concerned with numbers, like reasoning about continuous deformation of curved lines on curved surfaces, illustrated here. Perhaps our universe provides a wider variety of information mechanisms than the Turing-equivalent forms of computation can. (REF??? recent evidence in Biology.)

This project (since late 2014) uses the notion of construction kits and evolution of construction kits as a framework for understanding some of evolution's most complex achievements. As far as I know this is a new idea. (If not, please send me links to other work on evolution of construction kits.)

The components, relationships, and forms of composition of a construction kit (e.g. Meccano, Lego Bricks, plasticine, paper and scissors) together determine what entities can be constructed using the resources of the kit, how much they can vary.

What sort of construction kit had to be available from the earliest stages of this planet, to support not only all the physical forms and behaviours of life forms evolved on this planet, but also all the forms of information-processing, not only the information-processing involved in reproduction and growth, but also all the later forms including science, art, mathematics, engineering, politics, religious superstition, ethical debates, etc? Tibor Ganti (2003) proposed a minimal set of chemical mechanisms for elementary reproducing life forms. I don't know whether he thought his "Chemoton" idea sufficient to account for all forms biological information processing.

A more detailed discussion of construction-kits as explanations of biological possibilities was inserted here in November/December 2014: Sloman(C-K).

A closely related topic is the role of explanations of *possibilities* in science, discussed briefly here: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html.
A theory that explains possibilities may be deep science without being falsifiable.
(First steps toward a "generative grammar" for varieties of mechanism required for various forms of life and ecosystems, including information-processing mechanisms.)

CONTENTS

(8 Nov 2014) Entropy and evolution

Issues about entropy and the second law of thermodynamics also need to be discussed in detail eventually. I have a short, incomplete, comment in a separate document about entropy, chemistry, multi-stable, multi-structured dynamical systems and the *droguli* of Lionel Penrose:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html

KEY QUESTIONS

The Main question, discussed by many
How can a cloud of cosmic dust give birth to a planet full of living things as diverse as life on Earth?

Older questions
Many have asked: what sorts of physical and chemical mechanisms could make various stages of evolution possible, or various stages in individual development (epigenesis) in various types of organism, group or ecosystem.

They have also asked: what sorts of morphology (physical structure) and behaviour are needed at various stages of evolution or development, and what sorts of life-supporting chemistry are required. Examples of older questions, and links to further material, can be found online, e.g. at:
http://en.wikipedia.org/wiki/Hypothetical_types_of_biochemistry
http://en.wikipedia.org/wiki/Abiogenesis
http://en.wikipedia.org/wiki/Carbon-based_life
http://en.wikipedia.org/wiki/J._B._S._Haldane#Origin_of_life

Newer questions

This project asks:

-- What forms of information-processing (computation) and what information-processing mechanisms are required, to make the production and diversification of life forms possible?

-- What features were required in the Fundamental construction kit (FCK) provided by physics and chemistry before life began, that supported all the subsequent extensions and applications produced by natural selection, using many Derived Construction Kits (DCKs) (discussed in more detail in Sloman(C-K)).

-- How do the mechanisms, the forms of representation (encodings), and uses of information all evolve and develop and what new forms of life do they support, or in some cases interfere with?

-- What information contents were or are used by organisms, or parts of organisms, at various stages of evolution, at various stages of individual development, in various group interactions (mating-pairs, fighting pairs, parent-child, predator-prey, colony, culture, ecosystem, economy, ...)

In short, what are the causal roles of information in living things, and how do the information contents and the causal roles change over time, in individuals (at sub-cellular levels upwards), in species, in groups and in larger systems?

What information is, how many varieties there are, what can be done with it, what it can do, what mechanisms are required for these processes, are all complex questions discussed further below. An example: From dinosaurs to Birds (in another document).

Mathematical questions

What mathematical possibilities and necessities enable, constrain and shape the options for natural selection, for epigenesis, for individual competences, for cultures, for ecosystems?

What mathematical constraints? -- Topological, geometrical, physical, chemical, biological, computational, epistemological, linguistic, motivational?

D'Arcy Thompson, Brian Goodwin, and researchers included in a book of tributes to Goodwin, focused mainly on geometric and topological changes and constraints in evolution and development of physical forms (though I have not yet read all the papers carefully).

In contrast, the concerns of the M-M project include mathematical structures and constraints relevant to *types of information content, forms of representation of information, modes of reasoning, types of control of behaviour, forms of learning, and other uses of information* -- which are much less visible, and leave no fossil records. For these and other reasons, the study of M-M problems is still in its infancy. Far fewer researchers are equipped to think about these questions. At least physics, chemistry, and mathematics are taught to many children in schools.

Perhaps nobody is equipped yet: if some key ideas have not yet been discovered?

Other thinkers -- an incomplete list

If Alan Turing had lived longer he could have taken this project much further than I can. If Immanuel Kant had known what we know about information processing machinery, he would also have put the ideas to deep use. There are probably many other thinkers that I have not yet encountered whose ideas are relevant.

Some who have raised similar questions focus mainly on the evolution of human minds, e.g. Merlin Donald and Peter Gardenfors, among many others (including many I have not read).

There also seem to be overlaps with the work of Stuart Kauffman. Jack Birner (2009) has discussed ideas of Popper and Hayek related to this project.

Many of the writings of Daniel Dennett make points similar to the points made here, including his Kinds of Minds (1996), his writings on free will and others. However, we disagree regarding his claim of the centrality of "The intentional stance" for reference to mental states and processes, and his denial of the existence of the entities variously referred to as "sense-data" or "qualia", which I have argued arise naturally in sufficiently complex virtual machine architectures (e.g. in Sloman and Chrisley 2003 and the discussion of Virtual Machine Functionalism(VMF) above).

Daniel Dennett and I also seem to disagree on the origins of language -- which I argue must first have evolved for internal uses (with structural variability and compositional semantics, though not necessarily a "linear format"), to meet requirements for information processing within organisms, not communication between organisms.

> **http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111**
> **What are the functions of vision? How did human language evolve?**
> **Animals, including our pre-verbal ancestors, need "internal languages" (internal representations/encodings of information in the sense of Jane Austen below not Claude Shannon) for purposes that are not normally thought of as linguistic, including: perceiving spatial structures, having structured intentions, forming internal questions, working out what to do, initiating and controlling actions, learning things about the environment (including other agents), remembering, imagining, explaining, and planning. (That is not intended to be a complete list, merely an illustrative list.)**
>
> **These uses of information all require structured internal languages. Without internal languages learning and solving problems would be impossible. There would be no need for communicative languages if individuals had nothing to communicate, and had no internal means of storing and using information communicated or acquired by perception or learning. So, having one or more internal languages is a prerequisite for intelligent action and also for use of an external language Sloman(1979). Internal languages (forms of representation) must have evolved first, and must develop first in individuals: later both can develop in parallel. (It seems that the mathematician David Mumford, is one of the very few people who find this obvious Mumford(2016). This requires a "generalised" notion of a language: a GL. Internal GLs and external languages (ELs) require forms of representation that are manipulable, with**
> **- structural variability,**

- varying complexity (e.g. for information about objects/events of varying complexity),
- compositional semantics (allowing new meanings to be assembled)
See also: Biologically Evolved Forms of Compositionality (Added Sep 2018)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html

David Deutsch is pursuing partly related ideas referred to above, especially in his 2011 book:
The Beginning of Infinity: Explanations That Transform the World
His more recent "constructor theory" overlaps partly with the theory of construction kits under development here, but seems mainly to be focused on the bottom level: the fundamental construction kit, whereas this project emphasises multiple layers of derived construction kits mentioned above.

Terrence Deacon's 2011 book *Incomplete Nature: How Mind Emerged from Matter* overlaps with the ideas presented here in what seems to me to be a fairly shallow way, insofar as his book does not do justice to what we have learnt about information processing in the last seven decades, including work in Artificial Intelligence, on perception, reasoning, theorem-proving, language understanding, planning, vision, learning, and other topics related to this project. He rightly identifies the Second Law of thermodynamics as apparently contradicting the possibility of evolution producing increasingly complex stable forms of life, but does not seem to have encountered, or has not understood, the explanation based on quantum mechanics in Erwin Schrödinger's (1944) book. Although I have not yet read all of Deacon's book, like most contemporary philosophers he also seems to have no understanding of the achievements of designers of new kinds of multi-functional virtual machinery, illustrated in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html
Nevertheless, the book includes many interesting remarks about what needs to be explained.

Philip Warren Anderson "More is different"
Added 15 Feb 2015:

Iain Styles drew my attention to a very short but influential paper by the Nobel-prize winning physicist Philip Warren Anderson "More is different" in *Science* 1972 available on JSTOR, suggesting that there are many levels of organization between sub-atomic physics and the phenomena studied in other sciences, including biology and the social sciences. The suggestion that there are layers of construction kits of different sorts (proposed in Sloman(C-K)) appears to be supported by, or at least consistent with, his ideas.

Added 16 Mar 2015:
A separate document provides some notes on

Evelyn Fox Keller
Organisms, Machines, and Thunderstorms:
A History of Self-Organization
Historical Studies in the Natural Sciences,
Vol. 38, No. 1 (Winter 2008), pp. 45-75 and Vol. 39, No. 1 (Winter 2009), pp. 1-31
Summarised here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/keller-org.html

It surveys literature attempting to use ideas about dynamical systems to explain the emergence of mind, and argues that those ideas are inadequate to the task (a conclusion with which I agree, as explained in the paper on the roles of construction-kits in evolution).

Blind theorem-proving

Is evolution more like a blind theorem prover than a blind watchmaker -- proving theorems about what is possible?

Every time some new physical feature, behaviour, or mechanism arises in a living organism, that constitutes an implicit discovery that that sort of thing is possible, and was possible previously, though the realisation of the possibility may be more or less accessible at different stages of evolution. The evolutionary or developmental history contains an implicit proof that it is possible, but extracting the proof at the right level of abstraction may require sophisticated mathematical abilities that do not evolve till much later.

The meta-cognitive ability to notice that such a discovery has been made did not evolve till very recently (resulting from a mixture of biological and cultural evolution, among other things). It requires a highly specialised form of information processing competence, apparently available only to humans, among species evolved so far.

Yet evolution *itself* seems to have noticed some of them "implicitly", insofar as it discovered not only very particular solutions, but also generalised patterns that were then instantiated in diverse particular cases. The "laws of form" (studied by D'Arcy Thompson and others) illustrate this: A genome does not specify the precise shape and size of an organism or its parts, but rather a network of relationships between possibilities that can vary between individuals, but even more remarkably, can vary within each individual during that individual's growth and behavioural development (e.g. learning to control movements while size, shape, weight, weight distribution, needs and opportunities all change).

Another example of evolution discovering and using a collection of powerful mathematical abstractions is use of a basic collection of learning abilities to bootstrap abilities to learn how to use increasingly sophisticated features of the prevailing language or languages: a system that was eventually able to work in several thousand different cultures using different languages.

Moreover, the evolved mechanisms in humans somehow provide transitions between having various competences (possibly recently acquired) and becoming able to think about those competences and help others acquire them.

Those transitions from competences to meta-competences (using late developing genetic mechanisms, or learning), include the processes labelled "Representational Redescription" in Karmiloff-Smith(1992). They also have much in common with mathematical discovery, insofar as both often involve finding new abstractions that have many instances not previously encountered. (The recent claim that mathematics involves discovery and use of metaphors ignores the greater power provided by abstraction.)

Some of the ideas summarised in the above paragraphs are developed in connection with the Meta-Configured Genome (MCG) theory, illustrated in Figure Meta-Config, below, and explained in more detail with the aid of a video recording here: http://www.cs.bham.ac.uk/research/projects/cogaff/movies/meta-config/ and expanded in documents referenced there.

Discussion of the blind theorem prover idea continues below.

Evolving designs using blind mathematical composition

How do products of evolution combine with one another and with other environmental factors to form *niches* (sets of requirements) enabling and constraining future products of evolution (future designs partially matching the requirements) in multi-level dynamical systems constantly generating new dynamical systems, with new possible trajectories, and new feedback control mechanisms, in individuals, in social groups, in ecosystems, and now in multiple global villages?

How can the genotype available to a newly born or hatched animal make possible hugely (infinitely?) varied developmental trajectories in different environments, e.g. squirrels in different gardens with (mostly) shared genomes learning to defeat new "squirrel-proof" bird-feeders, and humans learning any one (or more) of several thousand very different human languages, absorbing whatever culture the child grows up in, acquiring competences relevant to local geographical features, local fauna and flora, local sources of food, shelter and danger, personalities of local conspecifics, etc. and in some cases creatively extending those environments through new inventions, new discoveries, new works of art, new moral teachings, new mathematical proofs, etc.

> **This is why proposing a behavioural test for intelligence is misguided (as Turing understood): no bounded behavioural test can establish the presence of all that potential. A test that indicates lack of intelligence may simply have been unsuited to the individual's capabilities. See http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-test-2014.html**

One common answer is that anything with human-like intelligence must use the same (postulated) general purpose learning mechanism, e.g. Juergen Schmidhuber, (2014). [The "empty mind" hypothesis mentioned below]

Even Turing (who should have known better?) toyed with that answer in his 1950 paper, though mainly in the context of a machine that learns to have text-based interactions. All the general purpose mechanisms I've heard proposed so far operate on compressing bit-strings, or symbol-streams, and don't seem to be capable of learning geometrical or topological facts or skills, including the competences of a squirrel, a weaver bird, or a mathematician studying properties of toroidal surfaces.

An attempt to characterise the sort of rich "Evo-Devo" interaction that makes nonsense of many speculations about evolutionary or environmental determination, led to the multi-stage developmental model of Chappell and Sloman (2007) depicted below -- extending Waddington's "epigenetic landscape" idea. (This sort of system must have been the product of a complex evolutionary trajectory: it could not have existed when life first began.)

Do we know enough about information and computation?

Are known forms of computation rich enough to provide such a genotype, or are there still secrets to be uncovered in products of evolution?

What are the (mathematical) properties of physics and chemistry that enable a protoplanetary dust cloud to produce machines that can ask questions like these? (The construction-kit question again.)

Is there something about chemistry that we have not yet understood? Only with the properties of chemistry do we seem to combine three features necessary for life from its earliest stages onwards: energy storage and transformation, mechanical and sensory structures that can act on the environment when appropriate, and mechanisms for storing, using, copying, and transforming information Schrödinger (1944), Ganti (2003).

Chemistry builds brains, at least in their early stages of construction, though chemical mechanisms remain essential for many brain processes throughout life. Perhaps interacting molecules do much more than we know, long after they have constructed neural mechanisms, as some neuroscientists now suspect, e.g. Trettenbrein(2016).

Related questions:
Have evolutionary and developmental processes produced biological machines that are intelligent enough to find the answers to these questions, or understand them if found? How?

Can schools and universities provide the sort of education required for researchers and teachers in this project?

CONTENTS

What we've learnt since 1952 - varieties of virtual machinery

The Meta-Morphogenesis project attempts to combine and extend Turing's ideas about morphogenesis and his earlier ideas about discrete computation, in the light of what we've learnt since 1952 from computer science, artificial intelligence, computer systems engineering, biology, neuroscience, linguistics, psychology, chemistry, physics, mathematics, and philosophy.

Unfortunately, because of flaws in our current educational systems, most philosophers and scientists seem to be unaware of the deep significance of what we have learnt about many forms of virtual machinery.

NOTE (24 Aug 2014):
Some of the references have been moved to a separate file, which includes documents on this web site relevant to the M-M project:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/m-m-related.html
Publications and references related to the Meta-Morphogenesis Project

The main focus of this project: Transitions in biological information-processing

A vast amount of research has been and is being done on the production by natural selection of new physical and geometrical forms of organisms, of many sizes and types, and production of new behaviours (e.g. J. Maynard Smith and E. Szathmáry, (1995), (1999), and Pallen (2009), among many others mentioned on Gert Korthof's web site).

The Meta-Morphogenesis (M-M or MM) project focuses instead on production of new types of *biological information processing*, including information-based control mechanisms, whether used for reproduction, growth, development, metabolism, perception, motor control, learning (including creation of new ontologies and new forms of representation), motive formation, planning, planned or unplanned behaviours, meta-cognition, communication, daydreaming, explaining, theory change, mathematical discovery, mathematical proofs, enjoying and producing art, or anything else. All new forms of computation that arise during evolution, development or interaction with other organisms are included. This requires use of a very general notion of "computation", or "information processing", that is not restricted to use of bit-based computers.

The changes in information processing include (a) *what* is done (as indicated in the previous paragraph), (b) why it is done, e.g. what benefits, if any, result, (c) what the information used is about (e.g. what it refers to, which can include past, present, future, remote, and non-existent entities, events, etc.) and (d) *how* all that is done, which refers to types of: information bearer, mechanisms for analysing, transforming, constructing, comparing, storing, retrieving information bearers, types of information processing architecture, combining different forms of information processing in larger wholes, types of self-monitoring, self-modulation, self-repair, self-extension, types of competition, types of conflict resolution, types of interrupt mechanism, use of virtual machinery, including multi-layer machines, distributed information-processing (involving several different individuals, or a whole community) and many more.

As explained below, the ability of natural selection to be a sort of "blind mathematician", discovering and using mathematical structures, seems to be crucial -- refuting philosophical claims that mathematics is a human creation.

CONTENTS

Why Meta-Morphogenesis?

.... Because the changes produced by the mechanisms of development and change include modified mechanisms for producing new changes in the mechanisms producing development and change.

Natural selection (or the biosphere) is a bit like a young child that has begun to learn, but has no idea that it is learning, what it is learning, how it is learning, why it is learning, what it will do with what it has learnt, why what it has learnt works and why what it has learnt sometimes proves inadequate, either for individuals or for whole species.

A difference is that over billions of years natural selection modifies its information-processing abilities far more than any child can do in a human lifetime, and modifies them in parallel in different ways. Moreover, at least on one planet, it has recently produced some individuals that have begun to understand some of what the evolutionary mechanisms produce without understanding.

The reproductive mechanisms do not normally produce ready-made full understanders, but individuals empowered to grow their understanding guided by the environment and by what some of their forebears and peers have already understood.

Some of those evolutionary changes bear a high level resemblance to processes in individual development in animals described as "Representational Redescription" in Karmiloff-Smith(1992). In particular, it seems that increases in competence both in evolution and in individual development involve mechanisms that partition discoveries into domains with mathematical structures that can be discovered by appropriate domain-related mechanisms (not merely the use of universally applicable statistical learning techniques as some have supposed). See also the quote from McCarthy below, and The Meta-Configured Genome discussed below.

A key idea: forms of creativity in evolution

Added 8 Aug 2014; Modified 16 Mar 2015

The key idea: evolution changes evolutionary processes and mechanisms, development changes developmental processes and mechanisms, individual learning changes individual learning processes and mechanisms, cultural evolution changes cultural evolutionary processes and mechanisms. Each of those is an example of meta-morphogenesis.

Moreover, each of these processes and mechanisms of change can impact on the others, over appropriate time-scales. That includes changing what evolution can do, by changing the resources available to natural selection -- e.g. by creating, modifying, and combining construction-kits available for evolution, development, learning and social/cultural processes Sloman(C-K).

If all that is correct, attempts to characterise any of those processes or mechanisms in a uniform way will lead to erroneous theories, because:

- **evolution is not a uniform process**
- **development is not a uniform process**
- **learning is not a uniform process**
- **social/cultural change is not a uniform process.**

Natural selection may seem to be a uniform process, but what it does depends both on the mechanisms generating options between which selections can be made, and the selection mechanisms, which in turn depend partly on external constraints and opportunities -- niches. The points summarised above imply that both the types of option and the selection mechanisms can change dramatically.

Those modifications include:

- changes in physical and chemical structures and processes (that require, and also make possible, more complex information processing),
- changes in reproductive machinery,
- changes in genome-driven or partly genome-driven patterns of individual development (epigenesis), that occur both across generations and within an individual's development,
- changes in the relative contributions of genome and environment and the stages at which they interact in individual development,
- changes in forms of adaptation and learning by individuals,
- changes in types of information-content that can be acquired, stored, manipulated, derived, and used; and changes in how particular sorts of information-contents can be represented (in physical or virtual media),
- changes in forms of sensing, perceiving and acting,
- changes in modes of communication and control between subsystems in an organism,
- changes in information-processing architectures within which diverse subsystems can interact, communicate, cooperate, compete and develop,
- changes in modes of communication and control between organisms,
- changes in types of cooperative or symbiotic processing,
- changes in requirements for and forms of competition,
- changes in abilities to acquire and use information about oneself and about other individuals (requiring two different but related forms of meta-cognition),
- changes in how parents influence offspring in their learning and development,
- changes in how groups of individuals acquire, use and transmit information,
- changes in how societies and cultures interact, including interactions involving new technologies,
- changes in the ways in which the physical environment produces new challenges and opportunities for information-processing in organisms of different kinds, including humans (sometimes as a result of biological processes, or as a result of other processes, e.g. geological events, asteroid impacts, climate changes, etc.) and
- Changes in Forms of consciousness, where consciousness is analysed as a far more complex concept than normally recognized because it exhibits parametric polymorphism (as I think Gilbert Ryle partly understood in 1949). This idea is under development in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html
- changes in the ways all these processes influence one another.

One of the most important discoveries of biological evolution was the power of "generative" forms of representation of information: e.g. encoding information using trees and networks of information, whose nodes can be either arbitrary non-decomposable objects, or structured (decomposable) objects composed of other objects, for example trees and networks whose nodes contain trees and networks.

The need for such meaning structures is clear in connection with the contents of complex sentences, pictures and diagrams, with parts that have parts that have parts. The need also exists in percepts, in mathematical formulae and proofs, in complex intentions, in explanatory theories, and in action plans.

The ability to create and operate on such structures has been a pervasive feature of AI programming languages, often described as symbolic programming languages, which typically also provide standard instructions for operating on numbers of various sorts. Without this sort of capability, human language, and, I suspect, powerful animal vision systems, could not have evolved. This is why widely used forms of representation using vectors of scalar values are not sufficient for explaining how organisms work. (They don't even suffice for representing chemical structures and processes.)

This is not intended to be a complete list of information processing novelties produced by natural selection. Extending the list, filling in details, and testing ideas by empirical research into processes and products of evolution, building working models to check the feasibility of the theories, and addressing a variety of closely related philosophical problems, including problems about relations between mind and body, are all among the long term aims of the M-M project -- potentially a huge, long term project.

Achieving such goals will require, among other things, major advances in AI and robotics in order to be able to test theories of how organisms work, and may even require novel forms of physical computing machinery, for instance if some of the functions of chemical information processing, with their mixtures of continuous and discrete changes, cannot be replicated in digital computers; and new kinds of mathematics may be required, for reasoning about how some of the systems work.

In the process we can expect many old philosophical problems to be solved or dissolved and many new ones to emerge.

The remainder of this document expands on some of these points and provides links to other, related documents on this web site and to relevant publications. (A partial list)

More on creativity in evolution/natural selection

(These ideas will be developed further in a separate document: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/creativity.html)

Although the processes start off "blind", the achievements are of a kind that would require highly creative processes of design, implementation, testing, development, debugging, and re-design, if produced by human engineers. In some of the later stages, when animal cognition begins to play a role in evolution, this is a form of conscious, but not yet self-conscious, creativity. Similar remarks can be made about varieties of creativity in development of individuals, discussed further in connection with "toddler theorems". (Compare Margaret Boden on creativity.)

A growing list of transitions in types of biological information-processing: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evolution-info-transitions.html
Biology, Mathematics, Philosophy, and Evolution of Information Processing

Mathematics is at root a biological, not an anthropological, phenomenon (as suggested by Wittgenstein). But its possibility depends on deep features of the universe, some of which evolution had to 'discover':
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html
An attempt to identify a major type of mathematical reasoning with precursors in perception and reasoning about affordances, not yet replicated in AI systems:
http://tinyurl.com/CogMisc/triangle-theorem.html

Even in microbes
I suspect there's much still to be learnt about the varying challenges and opportunities faced by microbes at various stages in their evolution, including new challenges produced by environmental changes and new opportunities (e.g. for control) produced by previous evolved features and competences -- and the mechanisms that evolved in response to those challenges and opportunities.

Example: which organisms were first able to learn about an enduring spatial configuration of resources, obstacles and dangers, only a tiny fragment of which can be sensed at any one time? What changes occurred to meet that need?

- Use of "external memories" (e.g. stigmergy)
- Use of "internal memories" (various kinds of "cognitive maps")

More examples to be collected here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evolution-info-transitions.html

---

CONTENTS

Are babies born with empty minds plus a learning machine?

Some researchers, including (as I understand him) Juergen Schmidhuber, (2014) seem to regard the pinnacle of evolutionary design as a totally general, domain-independent learning mechanism, which allows individuals to learn in any environment by discovering statistical relationships between sensory inputs and motor outputs; whereas there seems to be plenty of evidence that humans have different kinds of learning capabilities, used at different stages of development or for different domains of structures and processes.

Neisser and McCarthy on evolution and innateness

Ulric Neisser wrote in *Cognition and Reality,* W.H. Freeman, 1976.

> "... we may have been lavishing too much effort on hypothetical models of the mind and not enough on analyzing the environment that the mind has been shaped to meet."

John McCarthy wrote:

> *"Evolution solved a different problem than that of starting a baby with no a priori assumptions."*

"Animal behavior, including human intelligence, evolved to survive and succeed in this complex, partially observable and very slightly controllable world. The main features of this world have existed for several billion years and should not have to be learned anew by each person or animal." McCarthy (1996/2008)

McCarthy's suggestion is consistent with the hypothesis that natural selection produced a variety of different learning mechanisms useful for different stages of development in complex organisms, as depicted below.

One way to make progress on such questions is to try to chart the variety of forms of development of information processing in young animals including humans. A subset of that task forms the investigation into "toddler theorems" (the abilities of pre-school children to make proto-mathematical discoveries, without necessarily being aware of what's happening), described in a separate file:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html

NOTE: For a messy, still growing, collection of examples relating to learning and development see this web page on "Toddler theorems":
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html
(including an introduction to the idea of a "Domain").

The link with Alan Turing

(Apologies for repetition: will fix later.)
The idea of the Meta-Morphogenesis project arose from an invitation from Barry Cooper, co-editor of the award-winning book

"Alan Turing: His Work and Impact".
http://www.amazon.co.uk/dp/0123869803/
2013 PROSE Award announcements
Detailed list of contents and contributors.

After submitting my three promised papers I found that I was also expected to contribute to part IV (as a result of a misunderstanding). So I read Turing's 1952 paper on Morphogenesis Turing(1952), about which I previously had only very vague knowledge.

Turing's paper is not an easy read, especially for non-mathematicians. A useful overview for non-mathematicians by Philip Ball is Ball, 2015.

There is also a very readable introduction to the ideas in Margaret Boden's magnum opus Boden (2006). In particular, section 15.iv ("Turing's Biological Turn") gives a summary of Turing's work on chemistry-based morphogenesis (which she had read and admired decades earlier). The previous section of her book (15iii Mathematical Biology Begins) summarises relevant work by D'Arcy Thompson. E.g. she writes:

Accordingly, D'Arcy Thompson tried to relate morphology to physics, and to the dynamical processes involved in bodily growth. He suggested that very general physical (as opposed to specific chemical or genetic) constraints could interact to make some biological forms possible, or even necessary, while others are impossible.

Boden (2006) Vol 2, 15.iii.a: "Of growth and form" pp 1256

That is closely connected with the view of evolution as a "blind theorem prover", explained below.

Reading Turing's 1952 paper knowing that he had died two years after its publication, led me to wonder what he might have done if he had lived longer. My tentative (presumptious?) answer was that he might have worked on filling gaps in our understanding of evolution of biological information processing mechanisms of many kinds, used in reproduction, development, perception, control of actions, learning, problem solving, etc.

Newly evolved information processing mechanisms can alter the mechanisms of evolution (e.g. in mate selection) so I called this "The Meta-Morphogenesis Project".

The original proposal

The proposal for a Meta-Morphogenesis project, was first presented as a chapter (written in 2011), invited by Barry Cooper, as explained above, and published (in 2013) as part of the Turing Centenary volume:

A. Sloman, 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, 2013, pp. 849-856,
http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d

This was preceded by two scene-setting papers in the same volume:

A. Sloman, 2013, Virtual Machinery and Evolution of Mind (Part 1), in Eds. S. B. Cooper and J. van Leeuwen, *Alan Turing - His Work and Impact,* Elsevier, pp. 97--102,
http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106a
Discussing some features of virtual machinery that are potentially relevant to explaining the evolution of mind and consciousness, including their causal powers and differences between implementation and reduction.

A. Sloman, 2013, Virtual Machinery and Evolution of Mind (Part 2), in Eds. S. B. Cooper and J. van Leeuwen, Alan Turing - His Work and Impact, Elsevier, pp. 574--580,
http://www.cs.bham.ac.uk/research/projects/caogaff/11.html#1106b
Suggesting ways in which biological evolution may have taken advantage of virtual machines to produce self-monitoring, self-modifying, self-extending information-processing architectures, some of whose contents would have the defining features of qualia. This could provide a way for Darwin to answer the criticism that natural selection can produce only physical development, not mental states and consciousness.

A piece of evidence that Turing might have been interested in this topic: According to his mother, he had always been interested in living things, as depicted by her in this famous drawing:
https://www.commondreams.org/sites/commondreams.org/files/imce-images/turing_mother_drawing.jpg

Margaret Boden's commentary on Turing's work on morphogenesis provides this additional piece of evidence

> ***For the last few years of his life, Turing's energy went primarily into what he called "my mathematical theory of embryology". Indeed, after writing the first Manchester programming manual in 1950, he neglected his duties in the computing laboratory there as a result of his new interest.***
> **Boden (2006) section "15.iv. Turing's Biological Turn" (page 1261)**

Perhaps he would have moved (by analogy with some of his earlier moves) from studying embryology to studying the origins of embryology deep in the evolutionary past: the basis of this M-M project. (Later I'll discuss another link with Boden's work: her ideas on creativity and the varieties of creativity in natural selection, including ontological creativity, required for production of new types of virtual machinery mentioned briefly below. (Biological evolution is the single most creative entity that we know of: it directly or indirectly produced all the others.)

This is a complex, multi-faceted project, and could take several decades, or centuries (compare the history of physics, and chemistry).

Some of the main ideas are elaborated on this web page, and in other web pages referred to on a separate page. At present everything is provisional. The ideas have reorganised themselves several times since the first paper was written. Perhaps the most significant revision came from dawning realisation late in 2014 of the importance for evolution of fundamental and derived (including evolved) construction kits Sloman(C-K).

The Key Question: How could all life, and products of life, on Earth come out of a cloud of cosmic dust that converged to form a planet?

Steps towards an answer

Evolved information-processing -- in animals and machines.
(A huge, long-term multi-disciplinary project.)

Core questions and ideas

How can natural selection produce minds on a lifeless planet? A full understanding of our origins requires us to combine familiar ideas about natural selection with ideas unavailable to Darwin and Wallace, about evolution of information processing functions and mechanisms, since the simplest organisms in chemical soups billions of years ago.

Many research fields can contribute, including: genetics, microbiology, ethology, developmental psychology, neuroscience, linguistics, anthropology, philosophy of science, philosophy of mind, computer science, Artificial Intelligence and robotics. The research requires us to raise new questions about what evolution achieved and how it did so,

including questions about new forms of information, new uses for information, and new mechanisms for processing information. Doing that requires us to investigate new construction kits created, and then used, by processes of development, learning, and natural selection to support those developments.

Explanation by natural selection is not enough

As remarked above, Graham Bell wrote: "Living complexity cannot be explained except through selection and does not require any other category of explanation whatsoever." Like Bell, many writers on evolution (including philosophers) seem not to notice that adequate explanations need to mention both selection mechanisms and enabling mechanisms, as I suspect Bell is aware.

Without enabling mechanisms (referred to here as construction kits), selection processes will not have a supply of new working/viable options to choose from. In that case the selection mechanisms will no longer be able to select new viable options.

Both the selection mechanisms and the enabling mechanisms can change during evolution (partly by influencing each other). As a result, we can think of the initial enabling mechanisms, provided by physics and chemistry, as a form of construction kit that natural selection eventually uses (blindly) to build new mechanisms forming an enriched construction kit. If this happens repeatedly (as has happened spectacularly with computing mechanisms in the last 60 years or so), then the most recently evolved biological construction kits may be unrecognisable to scientists who know only about the initial mechanisms.

So the M-M project requires multi-disciplinary investigations of layers of evolved biological construction kits, some of which have helped to produce new construction kits for use by evolution.

There is a useful web site listing common misconceptions about evolution here:
http://evolution.berkeley.edu/evolibrary/misconceptions_teacherfaq.php
It does not bring out (or try to bring out) the full variety of types of explanation of evolutionary phenomena. E.g. Computer systems engineers have been discovering or inventing new types of information

processing for over half a century -- especially new types of virtual machinery. It appears that biological evolution made use of a similar discovery very much earlier, for good reasons, some of them summarised above.

CONTENTS

Relevant discoveries by biologists

Systems biologists are constantly discovering new biological types of informed control (information-based control). However, there may be types of biological enabling mechanisms (e.g. forms of chemical or biological computation) that we have not yet learnt about - and that may prevent us understanding some of the transitions in evolution, e.g. some changes in reasoning powers in our ancestors including changes from which we benefited.

Familiar ideas about natural selection need to be expanded to show how small changes can build up to create increasingly complex mechanisms involved in the processes that repeatedly produce:

a) new physical and chemical structures and processes supporting reproduction, metabolism, growth, immune responses, neural mechanisms, etc.;
b) new physical forms and new physical behaviours of organisms, including new types of sensing and acting;
c) New information-processing challenges, e.g. to deal with more complex physical phenomena, or more intelligent predators or prey, or to meet new demands on parents because of more sophisticated learning capabilities in offspring. (Challenges or requirements can evolve also, not only solutions. Challenges can come not only from new prey, new predators, new competitors, new physical environments, but also from new learning potential of offspring, or from new capabilities that are not easy to use.)
d) new information-processing capabilities and mechanisms, including sensory interpretation, motivation, learning, planning, decision making, interrupting, self-monitoring, teaching, etc.;
e) new evolutionary mechanisms, including new drivers of variation and new selection mechanisms.

Point (e) involves 'recursion': evolutionary morphogenesis changes mechanisms of evolutionary morphogenesis -- hence the label 'meta-morphogenesis'.

The Meta-Configured Genome
The project investigates how increasingly complex products of evolution produce increasingly complex forms of information processing including, new mechanisms of evolution -- generalising ideas in Turing's 1952 paper on chemical morphogenesis and also the theory of meta-configured individual cognitive development presented in Chappell and Sloman(2007), which includes an earlier version of this diagram, showing different times and different levels of abstraction at which information from the genome and from the environment combine (after varying developmental delays):

Fig Meta-Config

meta-configured
(Joint work with Jackie Chappell. Chris Miall helped with the original diagram.)

That theory (and diagram) referred to processes of development in an individual -- processes that change some of the mechanisms of later development in that individual.

The M-M project extends that idea to evolution, so that in this new context instead of the diagram referring only to development of individual organisms, it can also refer (loosely) to evolution of a species, or even of a whole ecosystem whose main features, including features affecting further evolution, change over time. Is there a corresponding diagram, relating the (metaphysical) creativity of the early physical universe to a genome supporting evolution of portions of the universe on various scales? ("Grounding" new metaphysical types -- Cf. Wilson(2017).)

This should be a strong constraint on explanatory cosmological theories, so far ignored by most physicists. Paul Davies is an exception, though I don't think his account includes the "multi-layered" genome idea above. Annette Karmiloff-Smith's ideas about "Representational Redescription" are also relevant.

The M-M project has begun to identify many changes in forms of biological information processing, including transitions in mechanisms of reproduction, mechanisms of learning and development, and inter-individual and inter-species forms of information-processing. Examples of distinct types of transition in biological information-processing are being collected here.

All this is inconsistent with claims that all natural forms of learning and discovery are special cases of some general type, e.g. probabilistic learning based on statistical evidence. (Which cannot account for discoveries of impossibilities and necessary connections, including mathematical discoveries, for example.)

An important under-studied transition is evolution of capabilities that led to proofs in Euclidean geometry long before modern mathematics, one of the most important extensions of human minds in the last few millennia. How did abilities to think philosophically evolve? Were the cognitive mechanisms unique to humans or did unnoticed subsets develop in other species? When will our robots begin to acquire these abilities?

The questions raised in the M-M project require long term multi-disciplinary collaborative research, perhaps comparable in scale to the Human Genome project. The relevance to philosophy of mathematics is discussed in a related web page.

9 Nov 2019: Some recent additions to the theory and discussion of MCG

"Information" -- a key idea for this project
(And for Jane Austen's novels!)

The concept of *"information used by organisms or machines or biological processes for various purposes"* is central to this project. But it is not the concept unfortunately labelled "information" by the great Claude Shannon and his many admirers. He understood the differences between the two concepts of information, but too many researchers ignore them. In fact many researchers think his is the only concept of "information" we have, or the only one relevant to science and engineering. But there is a much older one, used in everyday life and in science and engineering.

The older concept refers to information that has many causal roles, e.g. in evolution, in animal (including human) perception, learning, motivation, acting, interacting, thinking, asking, wondering, being puzzled, finding answers (etc.) Information of this sort is also used in processes of reproduction, though often in obscure, indirect ways.

This ancient concept was often used explicitly by Jane Austen over a century before Shannon's work, and by many others long before her. Several examples from her novel 'Pride and Prejudice' published in 1813, are presented here:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html
Jane Austen's concept of information (contrasted with Claude Shannon's).

However, I am not claiming that Jane Austen had considered all the uses of information relevant to biology. Readers may find it useful to try making a list of the kinds of information they use in a typical day, and what they use those kinds for -- or, more realistically, in a typical hour, such as the first hour after waking, including information used getting light (if needed), deciding whether to get up, getting out of bed, getting dressed, ...

In particular, "information-processing" here does not refer only to bit manipulation, or symbol manipulation, the operation of computers, or the sending and receiving of messages: those are all special sub-cases. The kind of information we are talking about does not need a sender and a receiver every time there is a use of information.

Acquiring information is finding out about something that the information refers to (or purports to refer to: it could be false information).

Information contents used by an organism can come from many different sources outside or inside the organism, recently or in the past, or both, and can play different roles: in controlling actions, in controlling sensing, in asking, answering or thinking about questions, in intentions, instructions, multi-step branching plans, conditions for doing something, predictive or explanatory theories, and many more.

All organisms, and many parts of organisms, including cells, use information -- and not just for reproduction. Information from external or internal sensors can turn on, turn off or modulate behaviour, which may