School of Computer Science | University of Birmingham UK

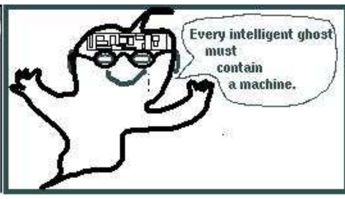

Every intelligent ghost must contain a machine

An information-processing machine

WARNING: DRAFT -- LIABLE TO CHANGE

(Comments and criticism welcome: Send to a.sloman[at]cs.bham.ac.uk)

Notes on Mathematics, Metaphysics, Evolution

Written after the final session of
[the Birmingham FramePhys conference]
Friday 11 Jan 2019
https://twitter.com/framephys?lang=en

Aaron Sloman
http://www.cs.bham.ac.uk/~axs/
Honorary Professor of Artificial Intelligence and Cognitive Science
School of Computer Science, University of Birmingham

Contents below

NOTE: A later document (April 2019) is intended to subsume and replace this one:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/emergent-physics.html (also PDF)

Installed: 15 Jan 2019
Last updated: 16 Jan 2019; 29 Jan 2019;30 Jan 2019; 9 Feb 2019; 18 Apr 2019

This paper is
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-framephys.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-framephys.pdf
It is an extension of the Meta-Morphogenesis project, pointing out some of the connections between products of biological evolution and varieties of dynamic metaphysical grounding.

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

Closely related:
-- Immanuel Kant's views on mathematics (1781)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html or pdf
-- Alan Turing (1938) on mathematical intuition vs mathematical ingenuity.
He thought only the latter could be implemented on digital computers Turing(1938) -- discussed in:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html or pdf
-- The multiple roles of compositionality in biology
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html (or pdf)
-- Jane Austen's concept of information (Not Claude Shannon's)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html (or pdf)
-- The Chemical Basis of Emergence in a Physical Universe (triggered by a later FramePhys workshop)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/emergent-physics.html (also PDF).
-- The Meta-Morphogenesis project
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html (Also pdf)

CONTENTS

Introduction: Two projects, a conference and Meta-Morphogenesis
Connections at a meeting of two projects
Questions about possibilities
Theme A. Mathematics, evolution, metaphysics
A hint from Turing
Another hint from Turing
Why chemistry?
A metaphysical exercise
Evolution's use of mathematics
Why are abstractions biologically useful?
From structures to processes manipulating structures
Evolution's repeated creation of new forms of information processing
Theme B. The timelessness of mathematics
(And spurious counterfactuals.)
No mathematical counterfactuals
Evolutionary boot-strapping and mathematical creativity
Unwitting mathematical competences
Evolved construction kits
Biological compositionality
A recent collection of ideas about evolution's (multiple) uses of compositionality.
Meta-morphogenesis: A related project:
Wilson on Metaphysical Grounding as Metaphysical Causation
Evolution's mathematical creativity
This Document
REFERENCES AND LINKS
APPENDIX 1. Historical interlude: recent developments in computing
APPENDIX 2. Background to the M-M project: evolution's metaphysical creativity
Creative Commons License

Introduction: Two projects, a conference and Meta-Morphogenesis

On 10-11 January 2019 the FraMEPhys project and the Metaphysical Explanation project hosted a conference on Metaphysical Explanation in Science, at the Ikon Gallery, Birmingham. The programme is here:
https://www.birmingham.ac.uk/schools/ptr/departments/philosophy/events/2019/framephys-conference-bham-1001.aspx

I was able to attend only the final session and, as expected, found connections, and disagreements, with some of my own work, the Meta-Morphogenesis (M-M) project, as explained below.

The FraMEPhys Project -- A Framework for Metaphysical Explanation in Physics -- is a five year research project, led by Alastair Wilson at the University of Birmingham. It aims to develop a new account of the contribution of metaphysics to how physics explains our world.
https://www.birmingham.ac.uk/schools/ptr/departments/philosophy/research/projects/framephys/index.aspx

An aspect of his proposal is construing grounding, a key notion in recent metaphysical debates, as a form of causation: *grounding = metaphysical causation* (or G=MC) Wilson(2018). I have been assembling evidence that various examples of metaphysical creativity evident in biological evolution can be understood as a form of dynamic metaphysical grounding, involving dynamic instantiation of timeless metaphysical types, including mathematical types. This is an example of metaphysical causation that is relevant to scientific understanding of many products of evolution.

In other words, evolution produces and uses instances of ever more complex mathematical structures in the designs it produces, i.e. structure-instances that instantiate timeless metaphysical types, including mathematical types and relationships. Some well known examples involving readily observable mathematical features of biological structures and processes were documented by Thompson(1917/1992), among others. Those are just the tip of an iceberg of mathematical creativity in the processes and products of evolution, including evolved forms of information processing.

Most biologists who, like D'Arcy Thompson, noticed some of evolution's uses of mathematical structures, apparently did not also notice or discuss (like Kant) the ability of some organisms (most obviously humans) to detect, and in some sense understand, necessary connections between aspects of those mathematical instances, e.g. the necessary connection between the number of sides, edges and vertices of a convex polyhedron. Such abilities require more than merely instantiation of the mathematical types. They require instantiation of mathematical reasoning capabilities, that currently are neither explained by neuroscience nor modelled in AI, although there are logic-based AI theorem provers that do something different, e.g. Shang-Ching Chou, *et al.* (1994).

I'll return to connections between the projects, and the link with mathematics, below.

The Metaphysical Explanation Project at the University of Gothenburg, led by Anna-Sofia Maurin, investigates the nature of metaphysical explanation https://metaphysicalexplanation.wordpress.com/. I suspect there are overlaps with the M-M project, but have not yet looked closely enough.

I was able to attend only the final session of the conference, on Friday afternoon 11th Jan. My main interest was in the last two talks:

14:55 - Atoosa Kasirzadeh (Toronto), "Can Mathematics Really Make a Difference?"
16:30 - Samuel Baron (Western Australia), "Counterfactual Scheming".

In discussing uses of mathematics in scientific explanations, they included material relevant to the M-M project, which refers both to mathematical aspects of the metaphysical content of biological evolution and to mathematical aspects of the products of evolution, especially the waves of increasingly complex information-processing mechanisms that enable, and extend, spatial reasoning competences, including the competences that led to ancient mathematical discoveries by Euclid and others.

A theme of the final session was discussion of mathematical counterfactuals, e.g. what would the consequences have been if 5 + 3 had been equal to 9, or if 17 had not been a prime number. I'll suggest below that counterfactual questions about whether a mathematical truth might have been a falsehood and the consequences thereof do not make sense, although questions about how what evolved might have been different make sense e.g. the ability to make mathematical discoveries, such as

5 + 3 = 8

exists, but it doesn't involve any change of truth-value of mathematical propositions only changes in mathematical competences and knowledge of evolved organisms.

Connections at a meeting of two projects

The remainder of this document expands on some connections between topics discussed at the conference and collections of ideas (inspired by Kant and Turing) that I've been working on, including ideas about explaining possibilities, i.e. attempting to answer Kantian questions of the form "How is X possible?" or "What makes X possible?" where X refers to some type of entity, process or state.

In particular I claim that the most important discoveries in science are discoveries of possibilities and explanations of possibilities, not discoveries of laws. Laws are of secondary importance, specifying restrictions of possibilities (i.e. impossibilities) as explained in Chapter 2 of Sloman(1978). A recent example was confirmation of the possibility of gravitational waves, predicted by the general theory of relativity. That now raises new questions, e.g. whether gravitational waves can be harnessed to transmit useful information.

QUESTIONS ABOUT POSSIBILITIES
(Updated 30 Jan 2019)
One of those questions is: what makes it possible for products of evolution to make and use deep mathematical discoveries? This is, in part, a question about evolvable kinds of minds and what makes them possible (grounds their possibility), and what made it possible for an ancient subset of such minds to make deep, still widely used, mathematical discoveries, including discoveries presented in Euclid's *Elements*, long before the development of modern, symbolic, logic-based, formal reasoning.

The answers may be different for different discoveries. For example, many animals, and pre-verbal human children, can discover and use possibilities that are expressed in terms of partial orderings, not precise metrics: e.g. it is possible for me to walk through some gaps in walls. For others it is impossible unless I walk sideways, though a young child may not need to turn sideways. Knowing such things does not require use of a metric for length (or distance), merely the ability to compare instances. A metric (up to a certain level of precision) can be developed later by adopting standard objects against which others can be compared, and then repeatedly subdividing those standard objects to provide standard objects with smaller differences between them.

I believe the answers to the "What makes it possible...?" questions are connected with evolution's production of spatial reasoning mechanisms in humans and many other species, including the deep spatial reasoning abilities of pre-verbal human toddlers. Both illustrate the ability to do mathematical reasoning without being aware of doing so, discussed briefly below.

These questions of the form "What makes X possible?" or "How is X possible?" seems to require answers that exemplify Wilson's ideas on grounding as metaphysical causation, referred to briefly below. There are also related questions about what grounds more widely spread spatial reasoning abilities in a variety of species, e.g. squirrels, nest building birds, apes, hunting mammals, and cetaceans as well as humans. (I've argued elsewhere that nothing in current (early 2019) AI or neuroscience, can explain or model those abilities.)

I append some notes related to the above and also to themes in the presentations I heard. If anyone has time/interest to spare, I'll welcome comments and criticisms. Feel free to copy to others who may be interested - or likely to criticise.

The next two sections (re-)introduce the main themes driving my comments. Further details come later, but can be ignored (or postponed).

The table of contents may help readers decide what to leave out!

(A) Main theme: linking mathematics, biological evolution, and metaphysics. (Expanding the above remarks.)

(B) Subsidiary theme: on the timelessness of mathematics, and resource limits.

B includes an objection to the idea of counterfactual conditionals in which the condition negates some mathematical fact. I raised this after Atoosa's talk, but I suspect it was also relevant to Sam's. (I did not keep detailed notes of reactions.)

After the summary of those two themes, the remainder of the paper discusses background and implications.

Theme A. Mathematics, evolution, metaphysics

In 2011 my thinking was disrupted by being asked to contribute to Elsevier's Turing centenary volume, including commenting on his paper on morphogenesis. For nearly half a century I had been trying to combine ideas from AI, computer science, mathematics, philosophy, psychology and biology, to expand and defend Kant's philosophy of mathematics -- the subject of my 1962 DPhil thesis, written before I knew anything about computers or AI. Reading the morphogenesis paper made me wonder whether Turing had been thinking about the possibility of new, chemistry-based (sub-neural in humans), forms of computation, combining discrete and continuous processes.

A hint from Turing
Part IV of the centenary volume was entitled: "The Mathematics of Emergence: The Mysteries of Morphogenesis". It included Turing's amazing, and very surprising, 1952 paper: "The chemical basis of morphogenesis", now his most cited paper (according to google) but generally ignored by philosophers and AI researchers. Thinking about it, and what Turing might have done later if he had not died in 1954, changed my way of thinking about philosophy, mathematics, AI and evolution.

> For non-mathematicians, Philip Ball has a very readable (and much shorter!) summary of Turing's paper:
> https://royalsocietypublishing.org/doi/full/10.1098/rstb.2014.0218
> Added 30 Jan 2019: There is also a video interview here:
> https://www.youtube.com/watch?v=6ed54_95kP4

Why did Turing write that paper on how complex patterns could emerge from fairly simple processes involving two substances spreading through a liquid yet forming patterns of varying complexity when they interact? His work on that paper may have overlapped in time with his work on the Mind 1950 paper, but there's only one (mysterious) sentence on chemistry in the 1950 paper, and no discussion of intelligence in the 1952 paper, whose final paragraph is:

> *"It must be admitted that the biological examples which it has been possible to give in the present paper are very limited. This can be ascribed quite simply to the fact that biological phenomena are usually very complicated. Taking this in combination with the relatively elementary mathematics used in this paper one could hardly expect to find that many observed biological phenomena would be covered. It is thought, however, that the imaginary biological systems which have been treated, and the principles which have been discussed, should be of some help in interpreting real biological forms."*

What would he have gone on to do in the next few decades, if he had not died two years later?

Another hint from Turing

There may be a connection with the distinction between mathematical ingenuity and mathematical intuition, in Turing's PhD thesis (published 1938), where he claimed that computers could replicate only mathematical ingenuity, not (human) mathematical intuition.

> I.e. he (implicitly) rejected what was later called the *generalised* Church-Turing thesis, and appears to have reached a conclusion about mathematical knowledge partly similar to Kant's claims about mathematics, as explained here:
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html (and PDF).

All of Euclid's axioms were important *discoveries* not arbitrary starting points to feed into theorem provers. He was, to some extent, retrospectively imposing a new structure (including the distinction between axioms/postulates and theorems) on a relatively unstructured collection of mathematical discoveries originally based on spatial reasoning, not logical reasoning. I don't know of any good mechanistic models or theories explaining how such spatial reasoning works.

There have been many automated (AI) geometry theorem provers based on *logical* forms of expression, and using modern logical forms of reasoning, many of them inspired by Hilbert's axiomatisation of the geometry presented in Euclid's *Elements* Hilbert(1899). I suspect Turing would argue that all those theorem provers illustrate what he meant by "ingenuity", especially symbol manipulating competences. As far as I know nobody has built (or knows how to build) an AI theorem prover that models the spatial reasoning (Turing's "mathematical intuition") leading to the original discoveries in geometry and topology, which some readers will have re-lived at school, though it now seems that in the UK, and many other countries, children are deprived of that experience, for bad educational, and sometimes bad meta-mathematical, reasons. (The Foreword by Robert Boyer in Shang-Ching Chou, *et al.* (1994) is very interesting in this context.)

Further evidence that no axiomatisation of Euclid's *Elements* can exhaustively account for our spatial mathematical knowledge are the geometrical discoveries that are not derivable from (e.g.) Hilbert's axiomatisation of geometry, such as the "neusis" construction, known to ancient mathematicians, which makes it easy to trisect an arbitrary angle -- which has been proved impossible in pure Euclidean geometry. (For anyone curious:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html also PDF)

Why chemistry?

The invitation to comment on the morphogenesis paper, plus my knowledge that after about 56 years, nothing in AI had so far modelled or replicated the ancient mathematical thinking leading up to and beyond Euclid’s Elements, left me wondering why Turing had turned from studying digital (discrete) computational processes to modelling continuous processes that can produce discrete spatial patterns.

A conjectured answer: he might have been thinking about ways in which chemistry-based (sub-neural) computing mechanisms, with their mixture of discrete and continuous processing (emphasised later by Schrödinger in (1944)), might fill gaps in our knowledge of how to model and explain ancient human mathematical intelligence (and also powerful spatial reasoning in many intelligent animals?).

That conjecture triggered a new project that I suggested Turing might have worked on had he lived several more decades, the "Meta-Morphogenesis" project: an attempt to characterise developments in biological *information processing* since the earliest (chemistry-based) stages of evolution, in particular developments that made it possible to start with

-- a metaphysically relatively sparse planet (or universe), containing only lifeless physical materials, structures, and processes

and eventually

-- without any external intervention, apart from steady supplies of energy and occasional disruptions by asteroids, climate changes, volcanic eruptions, tidal forces, etc.,:

-- populate a planet with an enormous and constantly changing variety of life forms, whose reproduction, development, and interaction with their environments depend on unwitting use of mathematical structures, including structures required for use of information processing mechanisms that, in turn, generate new useful mathematical structures instantiated in mechanisms of reproduction, growth, development, control, and interaction with diverse environments with differing mathematical structures,

-- thereby repeatedly adding new kinds of matter, structure, process, relationship, function, interaction, involving increasingly complex new metaphysical kinds, including needs, resources, growth, feeding, disposing of waste, various kinds of reproduction, controlling (pressure, temperature, chemical levels, location, orientation, relations to other things, ...), perception (of various kinds), recording/remembering, goals, goal formation, goal comparison, and many more!

Are there reasons for denying that these are metaphysical novelties, when they first occur?

Among the "many more" are kinds of perception, reasoning and action control, appropriate to interacting with 3D structures in spatial relationships with one another and with the perceivers/actors.

> Try walking through a forest, or a thickly planted garden, or a cluttered room, attending to all the changes in visibility of surfaces, relative distances, angles, portions of surfaces visible, angular sizes, and optical flow (texture flow), and other changing information fragments in different visible parts of the scene. As far as I know, there is nothing in current robotics, or current neuroscience, that explains how such information is processed and used (consciously or unconsciously) in brains of intelligent animals, including pre-verbal toddlers, as illustrated in this video:
> http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/small-pencil-vid.webm

There are many species with such capabilities but apparently only one developed an additional set of meta-cognitive capabilities enabling them to make discoveries about the spatial structures and relationships involved in their percepts and actions that eventually (through steps that are still unknown) led to the mathematical (especially geometrical) discoveries assembled by Euclid and many others.

A metaphysical exercise

Describing this (mostly still unknown) history of (some of) this planet involves in part a metaphysical exercise: characterising evolution's production of new *kinds* of entity (including new properties, relationships, structures, mechanisms, processes, capabilities, types of control, etc.) in which mathematical structures and relationships are involved.

Evolution's designs depend on mathematical constraints (necessities, impossibilities) as much as the designs of engineers and architects do.

Since the M-M project began, I've been trying to develop a multi-faceted theory of how, starting with the fundamental construction kit (FCK) provided by the physical universe and over time using increasingly complex derived construction kits (DCKs) of many kinds evolution might have produced increasingly complex life forms, with increasingly complex physical bodies, using increasingly complex forms of information processing in sensing and acting, then later in reflective thinking (e.g. mathematical thinking). If anyone is interested in any of these sub-topics, feel free to ask me questions or send me things you have written.

Evolution's use of mathematics

One of the themes that emerged in the M-M project was the importance of mathematics: evolution constantly "discovered and used" mathematical structures, including giving them causal roles in increasingly complex control systems, from control of motion through a gradient in a chemical soup to planning, designing, constructing, using, and maintaining increasingly complex products of human engineering. For example, there are multiple biological uses of negative feedback control (homeostatic control) of temperature, pressure, chemical concentrations, direction of movement, etc., long before the Watt governer was invented for controlling speed.

(These discoveries are sometimes misdescribed as uses of metaphors. Abstraction is more powerful and more general than use of metaphor, e.g. since it does not require constant reference back to particular exemplars, and it is more amenable to creative extensions -- a topic for another time.)

Why are abstractions biologically useful?

One reason for this increasing (implicit) use of newly discovered mathematical abstractions is that explicit encoding for all possible shapes sizes, forces to be used, and sensory configurations to be responded to in members of a particular species at various stages of development and in various "adult" contexts, would generate a combinatorially explosive amount of genetic material, and would require very much larger (exponentially larger, relative to the age and complexity of adults) search spaces to find the genetic configurations required. Genomes (and DNA sequences) would then have to be very much larger, and evolution would be drastically slowed down, by comparison with what has actually occurred. (This claim needs to be made more precise, and evidence assembled.)

This phenomenon can be compared with the differences between writing a computer program using machine code and writing a program with the same functionality in a high level, e.g. object-oriented, programming language. For simple designs machine code may be perfectly adequate (though not portable). For more complex designs the ability to use multiple previously designed abstractions that can be instantiated as needed can lead to much faster development and fewer design errors. It also simplifies modifying old programs to cope with new problems by making use of the function/parameter structure: old functions can be used with new parameters. In some cases dealing with novel problems leads to the discovery of new, more powerful, abstractions, with many new potential applications.

Software engineers (repeatedly) discovered all this in the twentieth century. I suggest that evolution made a similar discovery much earlier. (This needs a more precise argument, and assembling a lot of evidence.)

Finding powerful re-usable mathematical abstractions that can be deployed in special cases by supplying parameters, can change an intractable problem into a solvable one, and it seems that evolution did that many times. A very complex special case is the collection of abstractions (or more precisely, collection of collections of abstractions dealing with different developmental stages) acquired in human evolution that have already been able to generate several thousand different human languages (including spoken, written, and signed languages).

A structured organism with multiple body functions and body parts needs complex control mechanisms, for coordinated growth of internal and externally visible body-parts, for controlling temperature, fluid content, distribution of chemicals around the body, repair of damaged tissues, and for grasping, biting, walking, running etc., all continuing to function while the individual changes in size and body-mass distribution, and types of action performed. For all the structural details at all stages of development to be specified in a genetically inherited collection of parameters would be biologically impossible. (Perhaps I'll explain why in a later document, but it should be obvious.)

Using a collection of parametrised designs to solve the problem of specifying development and future possible behaviours requires evolution to be able to discover and use a collection of mathematical abstractions/structures, long before thinkers evolved who could make such discoveries, and reflect on them, teach them, use them etc.

The combination of biological evolution and the operations of its products, in generating increasingly complex novel structures and control mechanisms requires a vast amount of creative metaphysical causation, grounding increasingly complex and diverse forms of mathematically structured metaphysical novelty during evolution and use of biological designs, and in individual processes of perceiving, learning, and acting.

Having a large and growing collection of abstractions available allows increasingly rapid creation of new more complex abstractions, a phenomenon that has been demonstrated dramatically in the history of (metaphysically highly creative) human engineering. Moreover, software abstractions needed for controlling complex behaviours can be more rapidly recombined than hardware abstractions. I suspect evolution "discovered" that long before we did.

The power of mathematical abstractions in engineering contexts originally discovered and used by biological evolution was re-discovered much later by humans, which is why so much education of engineers now includes mathematics, and not just training in physical construction and manipulation (as may have sufficed for early makers of tools, clothing, dwellings, etc.).

In both biological evolution and recent engineering developments there has been a steady shift from increased physical complexity (e.g. larger cathedrals, bridges, cranes, etc.) to increased metaphysical complexity adding new "layers" of structure, function, process, causation, control and types of information used and produced.

From structures to processes manipulating structures
Until recently, most of the mathematics used by humans in engineering (including architecture) was concerned with specifying structures that met physical requirements of size, strength, durability, constrained and permitted relative motion of parts, etc. However, there has also been a steady increase in the use of machines that can replace humans in controlling processes that occur when changeable structures (i.e. mechanisms) are used. Examples include music boxes, piano-roll players, clocks of various sorts, self-orienting windmills, Jacquard looms for weaving cloth with different patterns, numeric calculators and punched card machines used for processing business data and controlling a wider range of machinery, long before the development of electronic computers.

All of that was a pale shadow of the mathematical sophistication in the discoveries made and used by biological evolution, without there being any intelligent agent observing, correcting, selecting, planning, and using the results. Many previously invisible biological examples have been made visible as a result of using high powered optical magnifiers and more recently very high speed cameras, e.g. to observe insect flight control processes, among many others.

If it were not for the many mathematical features and relationships constraining and supporting what is and is not possible for organisms, or parts of organisms (including microscopic and sub-microscopic parts of organisms) and controlling selection among possibilities (size, shape, speed, molecular structure, etc.) evolution could not have produced so many life forms with such diverse behaviours and diverse types of information processing.

A subset of human scientists and engineers now understand a lot more about what is possible in both natural and artificial control systems than was known a century ago, some of it documented in Boden(2006).

There have been past attempts to identify transitions produced by biological evolution Maynard Smith and Szathmary(1995). However, most of evolution’s transitions related to types of information, types of information processing, and types of information processing mechanism, were not included in their list of transitions, partly, I suspect (recalling interactions with MS when I was at Sussex University), because they were not well tuned-in to recent developments in computer science, software engineering and AI, especially symbolic, non-numerical, AI (e.g. planning, plan-debugging, theorem proving, problem-solving, and computational linguistics), which accelerated soon after the development of electronic digital computers. (Maynard Smith had been an aeronautical engineer before he became a biologist, so he had done a lot of numerical programming, which was extended in his simulated evolutionary experiments with changing numbers of predators and prey.)

Evolution’s repeated creation of new forms of information processing
One of the most important aspects of these developments was the constantly expanding need for new more complex forms of information processing, involving new types of information used in

more complex ways, including types of information with rich mathematical structures. The need for such developments in biological evolution can be compared with growing, increasingly complex, requirements for new kinds of information, based on new computational machines, languages, and tools, in human commerce, industry and engineering, including architectural developments going back thousands of years.

Compositionality is important in both contexts, biological evolution and human engineering (etc.), as discussed below in: Biological compositionality.

Attempting to document some of the most important examples of such transitions is a major goal of the M-M project. Unfortunately, I don't think the forms of information processing discovered so far by human scientists, engineers, and mathematicians are sufficiently rich and varied. Turing's comment that mathematical intuition (unlike mathematical ingenuity) is beyond the scope of computing machines can be construed as pointing out the need for research into novel forms of information processing, perhaps including the kinds of reasoning about continuous spatial structures and processes that led up to Euclid's *Elements*.

I have been collecting many examples and further developments of these ideas (all presented in online papers at various levels of incompleteness(!), including these:

Meta-Morphogenesis: Evolution of Information-Processing Machinery
http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d
The original proposal, published in the 2013 Elsevier collection:
Alan Turing - His Work and Impact

The current top level project overview (much expanded since the original):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html (also pdf)

Inconclusive discussion of Turing's comments on mathematical intuition mentioned above:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html

Theme B. The timelessness of mathematics (And spurious counterfactuals.)

This section is a reaction to questions raised during the conference about the possibility of counterfactual statements based on the possibility of some mathematical truths being falsehoods -- e.g. what difference would it have made if 17 had not been a prime number. I'll try to explain why this form of question is based on false assumptions about mathematics.

In what follows I'll refer to the mathematical creativity of biological evolution, insofar as evolution's products reflect mathematical discoveries that are used in increasingly complex products of evolution. This includes the development of perceptual and action control mechanisms that depend on new uses of mathematical properties of spatial structures and processes. But although we can ask what would have happened if a particular mathematical structure had not been discovered and used, we cannot sensibly asked what would have happened if some true mathematical proposition had been false.

The relevant mathematical structures (e.g. geometric, topological, arithmetical, and logical structures) are not products of evolution, and evolution's products cannot alter those structures, although the uses that are made of them can change.

A consequence of this is that it makes no sense to raise counterfactual questions about what might have happened if some mathematical truth had been false, e.g. if 5 + 3 had not been equal to 8, or if there had been a largest natural number or if the spatial intersection of a circle and a tetrahedron had been a cube.

No mathematical counterfactuals

I think it makes no sense to refer to the possibility of a mathematical truth being false, as happened more than once during presentations, though I don't recall exactly who said what -- Sorry!

There are no coherent mathematical counterfactual conditionals, though there are different mathematical structures (infinitely many of them, only a subset studied by humans, or instantiated in physical mechanisms, up to any time). The following two cases should therefore be distinguished:

B1: If a change is considered regarding a feature of mathematical structure M1, and that change is inconsistent with the remaining features of M1, then the proposed change is incoherent and cannot be used to identify a context for counterfactual or any other kind of reasoning.

B2: If a change is considered regarding a feature of mathematical structure M1, and that change is not inconsistent with the remaining features of M1, then the change will simply identify another structure M2, whose existence is not dependent on something *in M1* becoming false. However, insofar as such structures have complex combinations of features we can ask questions about which structures would be identified by various changes to the specifications, e.g. adding, removing or modifying axioms, provided no contradiction is generated.

For example, the parallel axiom in Euclidean geometry can be replaced with alternatives that define various non-Euclidean geometries, e.g. elliptical, hyperbolic and projective geometries. So asking what would be true in Euclidean geometry if the parallel axiom were false makes no sense. Instead we can ask: in which geometries are the parallel axiom false?

Likewise there are different number systems with different arithmetic properties, including: the natural numbers (characterised by Peano axioms), integers (positive and negative numbers), modulo arithmetic (integers modulo N, for each N), rationals, reals, transfinite ordinals, and various "strange" extensions, such as Abraham Robinson's non-standard arithmetic (which includes infinitely many numbers > 0 and smaller than EVERY rational number > 0. These "infinitesimals" can be (suggestively) modelled by "Horn shaped" angles between a circle and a tangent: as a circle grows larger the angle with the tangent grows smaller, but all such angles are smaller than every angle between two straight lines).

(The capitalised EVERY above prevents both rational and real numbers from satisfying Robinson's definition.)

Thanks to google, I've just learnt that C.S.Peirce had yet another notion of infinitesimal.

So mathematical structures are not the sorts of things about which we can discuss counterfactual questions. If structure M1 is changed in such a way to produce supposed structure M2, then different cases can be found, e.g.

-- If new features in M2 are inconsistent with the remaining unchanged features, then there is no mathematical structure M2 about which there are some new truths,

-- If there is no inconsistency (e.g. when a new axiom is added, consistent with all the original axioms in M1, or something is removed from M1 and replaced with a new axiom, without any contradiction), then the resulting structure M2 is a different structure from M1, and the discussion will be about a new structure not about the old one with a hypothesised new feature.

There are many examples in the history of mathematics where mathematicians at first failed to notice the existence of alternatives to the structures they knew about.

E.g. Imre Lakatos (1976) discussed the history of Euler's proof that:
    if a polyhedron has V vertices, E edges and F faces, then V - E + F = 2

Lakatos showed how discovery of various flaws in proofs led to discovery of a variety of previously unnoticed mathematical structures, and unnoticed subdivisions in old structures. E.g. does your concept of a polyhedron allow a polyhedron to have a triangular tunnel through it?

Evolutionary boot-strapping and mathematical creativity

During biological evolution, the machines in use at any time have essential roles in increasing the complexity and capabilities, especially the information processing capabilities, of their successors. Such mechanisms and their products have rich mathematical complexity. Different kinds of mathematics are relevant to different evolutionary layers. E.g. mathematical structures of syntactic forms used in human languages were causally irrelevant to microbes, insects, plants, and most other life forms for millions of years.

However, the *possibility* of those mechanisms existed in an inert, causally empty, way from the very beginning. Such possibilities could not have been realised micro-seconds after the big bang. Some possibilities have to "hang-around" (?) for a long time before they can do anything. Of course, that suggests that there are many possibilities whose time has not yet come, and perhaps many more (infinitely many?) whose time will never come. But the M-M project is concerned only with a tiny subset, of great importance to biology, science, engineering, and philosophy as we know them. At the core of that subset is a continually expanding variety of forms of information processing.

Among the realised possible information processing mechanisms, the M-M project seeks to identify especially key features of virtual machines running on the brain mechanisms produced by evolution (enhanced by other factors including cultures and various feedback loops) that made it possible for ancient humans, who knew nothing about modern logic, to make, discuss and document mathematical discoveries of great depth and power. Many of those ancient discoveries (e.g. concerning Pythagoras' theorem, and properties of prime numbers) are still in daily use world-wide, yet remain beyond the power/scope of current AI mechanisms, and beyond what current neuroscience can explain.

The original discoveries could not have used modern formal axiom systems, or the Cartesian representation of geometry in arithmetic. I suspect the mechanisms grew out of mechanisms for spatial reasoning shared with other species, including mechanisms for recognition and use of spatial **necessities** and **impossibilities**. But we may never know the exact history. Some examples of non-formal reasoning about impossibilities and kinds of necessity can be found in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html (also pdf).

Unwitting mathematical competences
Careful observations, e.g. by Konrad Lorenz, Jean Piaget, James Gibson, Philippe Rochat, Annette Karmiloff-Smith, observant parents (often helped by video cameras) and very many others, provide clues regarding closely related but unarticulated discoveries made by pre-verbal humans and other spatially intelligent species. (An example of a pre-verbal toddler making and using a non-trivial discovery in 3-D topology is available in this short video with commentary. What sort of language could her brain have used to represent the possibilities, her intentions, the constraints, the problems, and the solution? Does any neuroscientist have any ideas? Compare the spatial reasoning abilities of squirrels defeating "squirrel proof" bird-feeders in online videos.)

Different levels of evolutionary contribution to new metaphysical realisation

Clearly the spatial reasoning abilities in squirrels, humans and other apes are not found in microbes or insects, and once did not exist on this planet. At least two products of evolution need to be explained: the evolution of biological information-processing mechanisms that "blindly" create and use new control mechanisms (e.g. negative feedback control), and later the evolution of types of organism that can explicitly select some of those types of mechanism when they create new control mechanisms, e.g. humans designing self-orienting windmills.

An example of the first is evolution producing mechanisms for controlling temperature, osmotic pressure, motion (e.g. of roots) in the direction of increasing density of some needed chemical. An example of the second is evolution producing organisms (only humans so far?) that are also able to invent control mechanisms using negative feedback.

The first case involves implicit mathematical discovery by biological evolution in creating new species, and the second involves explicit mathematical discovery by individual organisms (humans, and possibly some other species) who think about and create mechanisms or select actions that are guided by negative feedback.

In the first case, the discoveries are implicitly encoded in genetic specifications and/or intermediate products of gene expression. In the second case the discoveries include more or less explicit information structures specifying what has been discovered -- as in a child or non-human animal noticing a solution to a practical problem, such as preceding a translation with a rotation.

A third level, that depends on evolution, or development, of additional reflective (meta-cognitive) capabilities is achieved when individual organisms become aware of what they are doing and why it is useful in some situations but not others. This requires not only solutions to be encoded but also analyses or explanations of properties of the solutions, encoded in thought processes, and documents, diagrams, and other devices used by humans as thinking aids or for communication.

These different abilities require different mechanisms providing different levels of sophistication -- a topic to be expanded elsewhere. Several AI researchers have worked on providing machines with related metacognitive debugging and learning mechanisms (one of the earliest being Sussman's program HACKER (1975)). But I don't know of any that come close to matching the spatial metacognitive abilities required for ancient mathematical discoveries, or a child's understanding of the fact that it is impossible to separate two solid linked rings (without any hidden slits), or to link two unlinked solid rings, without physical damage.

Yet more sophistication can occur in organisms that are able not only to reflect on their own successes and failures, but also to communicate what they have learnt to others, and engage in collaborative problem solving, debugging, ontology extension, etc. In some species the appearance of communication may be simply a genetically selected capability that is not understood by the individual with the capability, such as some of the unwitting teaching behaviours of carnivore parents feeding offspring.

There are different varieties of communication of acquired meta-knowledge across species, but as far as I know no other species comes close to the human achievements. Not all humans do, either because of individual genetic deficiencies, dysfunctional cultures, poor environmental support, or other forms of bad luck!

These are all cases where evolution directly or indirectly brings into existence instances of new types of mechanism with new types of capability. They are not merely temporally new, but also involve production of instances of types of information content that are not definable in terms of previously used types of information content. This requires "blind" creation by evolution of new more powerful forms of representation with associated mechanisms that implicitly assign new kinds of semantics to the representations.

The ability to use negative feedback is much simpler than, and requires simpler mechanisms than, the ability to manipulate and use information about causes and effects of negative feedback.

Likewise, temporally extended forms of meta-cognition producing new insights are required for the ability to reflect on and compare different cases of use, and to understand why some are successful but not others.

Yet another evolutionary change produces forms of communication between individuals that allow such discoveries to be communicated explicitly, thereby enormously speeding up and amplifying the effects of evolution. (Human teachers vary enormously in their abilities to communicate or install such competences in their pupils.)

So the M-M project addresses metaphysical questions about evolution's ability to produce new types of mechanism, with new types of power, acquiring and using new types of competence and knowledge. Those evolutionary discoveries of useful mathematical structures were replicated and extended more recently by (usually) professional (usually) highly educated mathematicians.

Many such humans extend mathematics as a largely introspective discipline driven by specially evolved motivational mechanisms, and not directly driven by practical needs, as discussed in Sloman(2009-2019).

These human discoveries build on but are different from evolution’s implicit mathematical discoveries shaped by deep practical (biological) engineering design problems, building on powers provided by the physical universe, then later building on powers produced by evolution itself. (An example of positive feedback.)

One of the best known examples of evolution’s powerful use of mathematical properties of the physical universe comes from Erwin Schrödinger’s little book on the chemical basis of life (1944), which shows how aspects of quantum physics, but not pre-quantum (e.g. Newtonian) physics, can explain how it is possible for chemical structures based on aperiodic mostly, but not completely, stable molecular chains to meet requirements for a reliable reproductive encoding mechanism. (It was published a few years before Shannon’s major publication and the discovery of the structure of DNA).

Our educational system does not normally produce thinkers able to move between different levels of analysis of function, type of mechanism and type of implementation, but this is a recurring requirement in the M-M project.

Their measures then violate the Archimedean axiom for arithmetic

for any N1 and N2, if N1 < N2 there is some natural number K such that KxN1 > N2

A great deal of mathematical discovery leading to new mathematical domains/structures has been triggered by "What if" questions. But these are not counterfactual questions as used in mathematical practice.

Evolved construction kits

The (still developing) theory of evolved *construction kits* including fundamental and derived construction kits of many sorts is discussed in Sloman (2014-18). These are not normally thought of as products of evolution, but they are essential for the production of the living things that are thought of as products, and biological construction kits (sometimes called toolkits) are produced during evolution of their products.

Biological compositionality
A recent collection of ideas about evolution’s (multiple) uses of compositionality.
Compositionality is often thought of only in the context of linguistic structure (e.g. in the Stanford Encyclopedia of philosophy (in 2018) https://plato.stanford.edu/entries/compositionality/, though its importance has also been recognized in computer science, in software designs and computer programs. What isn’t often noticed is that compositionality is also essential for evolving, developing, biological designs. That idea was triggered (in me) by having a short paper accepted for a conference on compositionality in September 2018 after which I expanded the paper in several directions, including a section on Metaphysical compositionality:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html (also pdf)
Biologically Evolved Forms of Compositionality
Structural relations and constraints vs Statistical correlations and probabilities

Although Wilson neither proposed, nor has since endorsed, these applications of his ideas, I think that (unless I am subject to wishful hallucination) his notion of *metaphysical causation* is naturally suited for use in evolutionary contexts (with appropriate extensions to match some of the details).

Moreover the evolutionary examples of *dynamic*, temporally extended, metaphysical grounding are partly analogous to his examples of grounding in a game of cricket, which can also be highly dynamic, with changing structures and relationships e.g. changing fortunes in a tense final innings. I can't recall whether he also pointed out the possibilities for further metaphysical creativity when a type of game interacts with a richer context, e.g. multi-game competitions, national pride, records being broken, and even (alas) spawning new kinds of criminality!

But a much larger variety of types of metaphysical creativity (including creating the possibility of mathematical minds) has occurred in our biosphere over millions of years than any philosopher has so far documented, although I think Kant (unwittingly) contributed key ideas to this project in his (often misunderstood) philosophy of mathematics. (Defended against standard criticisms in my 1962 DPhil thesis.)

Turing seems to have reached related, but less well developed, ideas in 1938 when he distinguished mathematical intuition from mathematical ingenuity and suggested (but did not argue) that computers can only achieve the latter. I've tried to tease out what he might have been thinking here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html (also pdf)

And finally all of this makes essential use of the notion of information, about which there are many theories, debates, confusions, and conflicting intuitions.

There are two central notions of information:
-- The fairly new, mainly syntactic, concept developed by Claude Shannon, to address a collection of technical (engineering) problems emerging from the telephonic/telegraphic services provided by Bell Labs
-- The much older concept of information as *semantic content* clearly understood and used by Jane Austen in her novels, written more than 100 years before Shannon's paper, documented in this paper using extracts from her novel "Pride and Prejudice":
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html (also pdf)

[I am not the first person to criticise misuses of Shannon's brilliant work (he wasn't confused about this) but I don't know if anyone else has invoked Jane Austen as a key witness. If so, please let me know.]

All this is a complex story with a lot of different interlocking strands. I don't really expect anyone to have time to look at all of it, but I'll be very pleased to receive criticisms and suggestions for improvement of any of it.

Meta-morphogenesis: A related project:

The Meta-Morphogenesis (M-M) Project is a combined philosophy+science, Kant- and Turing-inspired, project, using ideas from several scientific and philosophical disciplines, combined with mathematics and AI, aiming to explain how it is possible for a purely physical universe to produce, over time, a huge variety of life forms with increasingly sophisticated information processing capabilities, including some with abilities to make topological and geometrical discoveries like those in Euclid’s *Elements*. Some other species with considerable spatial intelligence may be using a previously evolved subset of those mechanisms.

The axioms and postulates listed by Euclid were not arbitrary (unjustified) assumed starting points for reasoning, as is permitted in modern formal systems. Rather they were products of deep evolutionary discoveries, probably built on evolved spatial reasoning capabilities shared with other intelligent species.

It is still not known what the spatial reasoning mechanisms are that humans and other animals deploy The mechanisms enabling humans to discover those axioms and postulates were products of complex, largely unknown, mechanisms that did not exist when this planet first formed, but were later brought into existence (caused to exist?) by a mixture of physical, chemical, and later biological mechanisms -- a metaphysical bootstrapping process that currently surpasses all other known, or human-designed, bootstrapping processes.
(For a brief survey of varieties of bootstrapping see https://en.wikipedia.org/wiki/Bootstrapping.)

Those evolved mechanisms and their products all included instances of abstract mathematical structures, that were discovered and used by evolution, but not created or modified by evolution. Mathematical structures cannot change, though different structures can be instantiated at different times and places, including different geometrical structures, such as Euclidean and non-Euclidean geometries, and different topologies.

The detailed operation of evolutionary mechanisms and patterns of reproduction were not discussed in the portion of the conference I attended, although there was some discussion of counterfactual mathematical premises in biological explanations mentioned below. However, the above remarks about mathematical structures imply that it makes no sense to ask counterfactual questions about such mathematical structures, such as "What would have happened if spatial containment had not been transitive, or if a certain number had not been a prime, or if a triangle had had four sides, or if the set of natural numbers had been finite?"

We can ask what would have happened if some portion of the universe, or some product of evolution, had instantiated a different mathematical structure, but not what would have happened if the structure had been different. E.g. we can ask what would have happened if the reproductive cycle of a certain species had not used a prime number of years, but we cannot ask what would have happened if 17 had not been a prime number, or if 17 and 19 had not been co-prime.

This is not because the mathematical truths are all trivial matters of definition -- "analytic" in Kant’s terminology -- since more than matters of definition are at stake. (I think Kant’s philosophy of mathematics was basically correct, but needed some clarification and qualification, e.g. to allow that mathematical discovery processes are not infallible, even if they lead to necessary truths. (For more on Kant’s view of mathematics see links above.)

More difficult questions in the M-M project are concerned with how information about mathematical structures came to be encoded in genomes and used in evolutionary designs, and how those encodings influence processes of individual development, and what the consequences for individual organisms are. A (tentative, still incomplete) answer (co-developed with Jackie Chappell) refers to a complex multi-layered process of gene-expression alluded to in the label "Meta-Configured Genome" Sloman&Chappell (2007-2018) used below. This involves quite different theories of learning and development from those used in current neuroscience and AI.

Wilson on Metaphysical Grounding as Metaphysical Causation

I have found it useful to express some aspects of the M-M theory using ideas borrowed from Alastair Wilson linking grounding and causation, e.g. in

"Metaphysical Causation", Nous 52(4):723-751, 2018.

Also available here:

http://alastairwilson.org/files/mcweb.pdf

Around 2014 I first heard him talk about *grounding* as *metaphysical causation*:

"Grounding = metaphysical causation" (G=MC)

and felt that that idea could usefully summarise ideas I had begun to develop about how biological evolution introduces astonishing novelty into a physical universe.

Metaphysics is often construed as a study of timeless truths, whereas Wilson discussed grounding in processes, as in a cricket match in which the rules of the game plus events on the field produce new states: i.e. Wilson's metaphysics is not timeless. This is not a claim that human metaphysical theories change, but that what kinds of things exist, and can exist in particular times and places can change, which calls for an explanation of what makes such metaphysical changes possible, in contrast with scientific explanations of how some specific type actually came to have instances.

> This can be contrasted with P.F. Strawson's ideas on "Descriptive Metaphysics", (1959). If I recall correctly, he defines descriptive metaphysics as the study of a *fixed* collection of beliefs common to all human beings at all times, which he contrasts with products of "Revisionary Metaphysics" which seeks to criticise and replace common ideas about what exists. I have some partly critical comments on Strawson's ideas here:
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-descriptive-metaphysics.html
> (This is written from memory of reading Strawson a long time ago. Corrections and criticisms welcome.)

In contrast with Strawson, Wilson explicitly discusses metaphysical processes of change, e.g. processes in a cricket match). These ideas can be extended to metaphysical changes of far greater significance (e.g. metaphysical changes in the processes of biological evolution). This is not a matter of "revisionary metaphysics" in which human metaphysical claims change. Rather Wilson acknowledges that there are metaphysical processes in which new kinds of entity come to exist -- unless I have misunderstood.

I have been using (not, I hope, mis-using!) his ideas in discussing the metaphysical creativity of biological evolution: the most creative mechanism known in our universe. Although mechanisms of biological evolution cannot change the variety of possibilities (possible organisms, behaviours, products of organisms, etc.) supported by the physical universe it can have a massive impact on the reachability some of them, or the time required to reach them. When the earliest life forms were (somehow) produced on this planet the variety of new life forms that were reachable by physical chemical processes within a century was minuscule compared with the variety of new life forms reachable within a century now, which includes not only new species capable of being produced by natural processes in ecologically rich portions of the planet, but also new species capable of being produced by intended or unintended human intervention.

Soon after the earliest life forms appeared, only very simple life forms could be reached in a short time, whereas now the variety of new forms that could be reached within (e.g. a century) includes new forms with many more degrees of complexity, including new plants and animals. This is partly because of the existence of many evolved biological construction kits.

Evolution constantly uses aspects of a range of possibilities supported by the universe at a particular time to extend the range of possibilities supported -- a positive feedback loop that enables the rate of evolutionary change to be substantially speeded up, using a steadily increasing supply of new construction kits, as discussed in Sloman (2014-18).

Insofar as biological information processing mechanism use chemical processes, in which there are both discrete changes (e.g. formation and unlocking of chemical bonds) and continuous changes, e.g. folding, twisting, moving together, moving apart, the biological machinery is not restricted to discrete processes as Turing machines and digital computers are, unlike so-called analog computers, most of which have been displaced by approximately equivalent but much cheaper and more reliable digital computers, where exact equivalence is not required. For examples and discussion see https://en.wikipedia.org/wiki/Analog_computer

Evolution’s mathematical creativity

Modified 9 Feb 2019

Biological evolution exhibits mathematical creativity, insofar as evolution produces and uses instances of increasingly complex mathematical structures and also produces increasingly complex mechanisms for discovering and using new mathematical structures and relationships. In that way, biological evolution has mostly been far in advance of human engineers. (I am not suggesting that Wilson agrees with any of this: however, these examples of evolution producing new types of organism, with new types of capability, seem to me to illustrate his notion of metaphysical causation, at least as well as his examples of metaphysical causation during a cricket match.)

N.B. I am not claiming that evolution produces any mathematical structures. Mathematical structures exist "timelessly". However, from time to time, evolution does produce new instances of mathematical structures that had not previously been instantiated in this universe (or on this planet).

It also produces new information processing mechanisms capable of discovering and using some of those mathematical structures. We can loosely think of an eternally available "cloud" of eternally enduring mathematical structures, from which evolution selects and instantiates increasingly complex specimens. That process continually extends evolution’s abilities to find and use new mathematical structures. (More precisely, it continually extends the collection of evolutionary mechanisms that have been instantiated and used.) Is this what Plato was trying to say about mathematics???

N.B. I hope my constant reference to what evolution does is not taken to imply that I regard biological evolution as a kind of conscious agent performing intentional actions. Evolutionary processes do wonderful things in the same sense of "do" as tornadoes do dreadful things. However, many products of evolution can perform intentional actions (e.g. mate selection) and that can greatly extend what evolution achieves.

Neither should it be assumed that all the information processing done by products of evolution is done entirely inside the organisms. As remarked in chapters 6 and 8 of Sloman(1978), there are many contexts in which humans and other intelligent animals use information stored and manipulated outside their bodies, e.g. in doing mathematics with the aid of diagrams, or aligning objects in order to compare sizes, angles or shapes. (Unfortunately this discovery has provoked some researchers into denying the existence or use of internal information processing!)

Eventually (and surprisingly?) evolution (assisted by its previous products and aspects of the physical world) even instantiates designs for mathematical minds that can make mathematical discoveries at far greater speeds than evolution itself can. It does this partly on the basis of mechanisms that select generic available design patterns during individual development on the basis of results of previously instantiated selections. (The most obvious example is linguistic development, also discussed by Karmiloff-Smith (1992).)

This activity of a "meta-configured" genome Sloman&Chappell(in progress) allows an evolved specification for a collection of sub-genome patterns to be instantiated step-wise, partially under the influence of the environment and its effects so far, a *far* more powerful form of individual development than any so far discovered in psychology, neuroscience or AI (partly identified by Annette Karmiloff-Smith).

Another illustration of this idea (also noted in the work of Noam Chomsky and other theoretical linguists) is the fact that language-generating aspects of the human genome can produce thousands of very different languages, that differ in far more complex ways than height, limb-length ratios, eye-colour, muscular strength, etc. Currently investigated "deep learning" neural mechanisms, using a multi-layered learning tower that exists from the start in each individual, lack the required mathematical creativity of a meta-configured genome which does not allow "higher levels of the genome" to be instantiated until results of earlier learning are available to influence the instantiation of the higher level abstractions, which then seek novel structures using previously instantiated patterns combined with new and old data. (This needs more detailed discussion. People familiar with research on language development and/or ideas of Karmiloff Smith may recognise some of what is being claimed.)

A partially related technique in AI is *Genetic Programming*, which replaces a linear inherited genome with an inherited collection of tree-structured patterns of varying complexity, where sub-trees represent more or less complex design fragments found useful in previous evolutionary discoveries, as summarised in
https://en.wikipedia.org/wiki/Genetic_programming

However, evolution’s use of chemical structures that can be combined in richer ways during epigenesis may have significant consequences not yet worked out.

This is part of a general research strategy of replacing hypothesised natural selection mechanisms that can only passively select between alternatives that happen to have turned up through random mutations, with more powerful mechanisms that can generate and combine new sub-types of previous types of mechanism far more rapidly during individual development. E.g. instances of designs for predators interact with instances of designs for their prey, enhancing already instantiated designs in members of both species.

That can be preceded by interactions between immature instances of the same design (through "play") in ways that help select instantiation options for a generic genome, in preparation for later interaction with "the real thing" (real prey or predator). The development of language in humans also uses such meta-configured genomes, but in an even more complex, multi-stage process, that doesn't involve risks of being eaten, etc. Sloman(2015).

This Document

This document presents some aspects of the M-M project that I think are closely related to the two projects represented at the conference, and could lead to extensions and/or corrections of some of the ideas emerging from those projects and related work presented at the conference, especially links between physics, biology, mathematics, and varieties of consciousness, just as attending the conference has helped me clarify the metaphysical presuppositions and implications of the M-M project. I welcome comments, criticisms and suggestions,

---

APPENDICES BELOW, AFTER REFERENCES

APPENDIX 1. Historical interlude: recent developments in computing
APPENDIX 2. Background to the M-M project: evolution's metaphysical creativity

---

REFERENCES AND LINKS

A. Sloman (2018-Kant) Key Aspects of Immanuel Kant's Philosophy of Mathematics ignored by most psychologists and neuroscientists studying mathematical competences.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html (and pdf)

Margaret A. Boden, (2006), *Mind As Machine: A history of Cognitive Science* (Vols 1--2)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/boden-mindasmachine.html

Shang-Ching Chou, Xiao-Shan Gao and Jing-Zhong Zhang (1994), *Machine Proofs In Geometry: Automated Production of Readable Proofs for Geometry Theorems,* World Scientific, Singapore,
http://www.mmrc.iss.ac.cn/~xgao/paper/book-area.pdf

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

J. J. Gibson, 1966, *The Senses Considered as Perceptual Systems,* Houghton Mifflin, Boston.

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

David Hilbert, (1899), *The Foundations of Geometry,*
Translated 1902 by E.J. Townsend, from 1899 German edition. Available at Project Gutenberg, Salt Lake City,
http://www.gutenberg.org/ebooks/17384

Immanuel Kant, 1781, *Critique of Pure Reason* Translated (1929) by Norman Kemp Smith, London, Macmillan.
http://archive.org/details/immanuelkantscri032379mbp

Annette Karmiloff-Smith, 1992,
Beyond Modularity,
A Developmental Perspective on Cognitive Science,
MIT Press (1992) -- My informal, unfinished, review:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html

Imre Lakatos, 1976, *Proofs and Refutations,* Cambridge University Press, Cambridge, UK,

George Lakoff and Rafael E. Nunez, 2000, *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*, Basic Books, New York, NY,

C. Maley and G. Piccinini, 2013, Get the Latest Upgrade: Functionalism 6.3.1. *Philosophia Scientiae,* 17(2), 1-15.
http://poincare.univ-nancy2.fr/PhilosophiaScientiae/

J. Maynard Smith and E. Szathmary, 1995, *The Major Transitions in Evolution* Oxford University Press, Oxford, England:,

Otto Mayr (1970) *The Origins Of Feedback Control*, MIT Press, ISBN 9780262130677
https://archive.org/details/TheOriginsOfFeedbackControlOttoMayr

Erwin Schrödinger (1944) *What is life?* CUP, Cambridge,
I have an annotated version of part of this book here
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html

A. Sloman [1978, revised], *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind*.
http://www.cs.bham.ac.uk/research/projects/cogaff/crp/

A. Sloman, 2009-2019, Architecture-Based Motivation vs Reward-Based Motivation, Originally published in *Newsletter on Philosophy and Computers,*
American Philosophical Association, 09, 1, pp. 10--13, Newark, DE,
Online version much revised and extended
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html

A. Sloman, (2013) Virtual Machine Functionalism (The only form of functionalism worth taking seriously in Philosophy of Mind and theories of Consciousness),
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html

A. Sloman (2015), Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)?, Online presentation, School of Computer Science, University of Birmingham, UK,
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang

A. Sloman (2016-creative), The Creative Universe (Draft discussion paper).
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/creativity.html

Aaron Sloman and Jackie Chappell (2007-2018) The Meta-Configured Genome (work in progress)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html (also pdf)

Aaron Sloman (2014-18), Construction kits for evolving life
[In progress. Begun Nov 2014]
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.pdf

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

Galen Strawson, 2019, A hundred years of consciousness: a long training in absurdity, *Estudios de Filosofia*, 59, pp. 9--43,
https://aprendeenlinea.udea.edu.co/revistas/index.php/estudios_de_filosofia/article/view/333363

Gerald J. Sussman, (1975) *A computational model of skill acquisition,* American Elsevier, San Francisco, CA,
http://dspace.mit.edu/handle/1721.1/6894

D’Arcy Wentworth Thompson (1917/1992) *On Growth and Form*.
The Complete Revised Edition (Dover Books on Biology) Originally published 1917.
http://www.amazon.com/On-Growth-Form-Complete-Revised/dp/0486671356/ref=cm_cr_pr_orig_subj

A. M. Turing (1938), (Published 1939) Systems of Logic Based on Ordinals, *Proc. London Mathematical Society*, pp. 161-228, https://doi.org/10.1112/plms/s2-45.1.161

A. M. Turing, (1950) Computing machinery and intelligence, *Mind*, 59, pp. 433--460, 1950, (reprinted in many collections, e.g. E.A. Feigenbaum and J. Feldman (eds) *Computers and Thought* McGraw-Hill, New York, 1963, 11--35),
WARNING: some of the online and published copies of this paper have errors, including claiming that computers will have 109 rather than 10^9 bits of memory. Anyone who blindly copies that error cannot be trusted as a commentator.

A. M. Turing, (1952), 'The Chemical Basis Of Morphogenesis', in
Phil. Trans. R. Soc. London B 237, 237, pp. 37--72.
(Also reprinted(with commentaries) in S. B. Cooper and J. van Leeuwen, EDs (2013)). A useful summary for non-mathematicians is
Philip Ball, 2015, Forging patterns and making waves from biology to geology: a commentary on Turing (1952) 'The chemical basis of morphogenesis', *Royal Society Philosophical Transactions B,* http://dx.doi.org/10.1098/rstb.2014.0218

Alastair Wilson(2018) "Metaphysical Causation", *Nous* 52(4):723-751
http://alastairwilson.org/files/mcweb.pdf

Appendix 1. Historical interlude: recent developments in computing

In the early days of electronic computing, most computer programming done by scientists and engineers (though not Turing) was numerical, using languages like Fortran, whereas the last 70 years or so has seen a steady increase in non-numerical forms of computation including manipulation of symbols, syntactic forms for complex symbolic structures, increasingly complex forms of computer programs for specifying and manipulating structures (including program-generating programs, i.e. compilers), management of networks of such systems (e.g. internet-based email, reservation systems, marketing systems, educational systems, game-playing programs, program-generating systems, and many more).

One of the early developments in that direction, inspired by the work of Grace Hopper, was development of COBOL an English-like language for business applications, while around that time symbolic programming languages for AI were being developed, e.g. LISP in the USA and POP2 in Edinburgh, followed by a flood of different languages soon after.

The recent (apparent) successes of "deep learning" techniques involved a shift in the reverse direction (back to number-crunching), but I expect its deep limitations will become increasingly evident in the next decade or so, e.g. because statistical reasoning cannot lead to conclusions about what is impossible, or necessarily the case, which underly ancient mathematical discoveries and many aspects of spatial reasoning in humans and other animals. (as Kant pointed out: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html .)

Evolution's symbolic specification languages (still mostly unknown) were far more powerful and made possible, among other things, the later development of thousands of (non-numerical) human languages for thinking, communicating, and reasoning, followed (later?) by processes of mathematical discovery leading up to Euclid's *Elements*, using information processing mechanisms that have not yet been replicated in AI (or explained by neuroscience).

NB No human language could be innate, not even French. Moreover, human languages cannot be acquired solely by learning from competent speakers, otherwise the process could never have started. Therefore evolution produced language creation mechanisms not language learning mechanisms. But the creation is collaborative and the majority will normally be ahead of newcomers. A stunning example was the **creation** of a new sign language by deaf children in Nicaragua, where there was no pre-existing community of experts.
A short video report is here (BBC?):
https://www.youtube.com/watch?v=pjtioIFuNf8

I know of no theory/model of learning that can explain or replicate such processes. The metaphysical creativity of AI researchers and cognitive scientists still lags far behind the creativity of evolution.

APPENDIX 2. Background to the M-M project: evolution's metaphysical creativity

Ancient humans made sophisticated mathematical discoveries (with properties partly identified by Kant, as discussed in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html (also pdf). Many of those mathematical structures had not been instantiated when life forms first emerged. The ancient discoveries in geometry, topology and number theory are still in use thousands of years later. Although mathematical structures and truths are timeless (as discussed below), the information-processing mechanisms that have been instantiated, and the mathematical structures used, can change, steadily increasing the information processing powers available to evolved species.

> NOTE: "Information" is here being used in the sense of Jane Austen, rather than Claude Shannon, as explained in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html also (pdf). Shannon information refers merely to properties of information vehicles, not information contents. There are complex relationships between them.

"Information-processing" here refers not to manipulation of bit-patterns, which are merely useful vehicles for information storage transmission and use in a subset of current artefacts produced by humans. Contrast the rich, rapidly changing, three-dimensional information used by a bird flying through branches and foliage to land safely on its nest, using spatial intelligence unmatched by anything in current AI, and unexplained by current neuroscience. Biological information processing is much older, much richer, and more varied in its forms, contents and mechanisms -- many not yet understood (e.g. sub-neural, chemistry-based, brain mechanisms). (James Gibson began to characterise some of this complexity in his work on perception of affordances, (1979), but barely scratched the surface of the topic.)

Biological information processing mechanisms include not only physical/chemical machinery but relatively recently evolved *virtual* machinery that is still barely understood by most philosophers and scientists studying minds and brains. A partial account of what’s missing is in Maley & Piccinini(2013), extended in Sloman(2013).

A feature of the M-M project is the increasingly complex roles of multiple layers of concurrently acting biological virtual machinery performing information-processing tasks that repeatedly produce new types of states and processes with new causal powers, involving new mathematical structures. Recently evolved virtual machines are so different from physical and chemical processes that they seem to many thinkers to be incapable of being explained or modelled in physical machines. The scathing attack on modern philosophers of mind by Galen Strawson (2019) seems to be based on a dim appreciation of some of the features of virtual machines, though he never mentions them or acknowledges that they can exist in complex running computer-based systems, e.g. the internet.

There’s a lot more to be said about unnoticed features, including metaphysical features, of virtual machines, ignored by most philosophers because they underestimate the complexity and indirectness of the relations between physical machines and the most complex virtual machines (e.g. the internet-based email virtual machine that now links minds, brains, phones, computers and a vast amount of networking technology. Maley & Piccinini(2013) go further than most in the right direction, but even they miss some of the important points that emerge during personal experience of designing, implementing, testing, debugging and extending complex virtual machines in which causation can go "upwards" from physical through virtual machinery, "sideways" within and between VMs and also "downwards", e.g. from decisions in virtual machines to changes in physical behaviours.

This is clearly a type of metaphysical novelty/creativity/causation as worthy of philosophical study as any other, and more important than most, because of its deep potential engagement with hard unsolved scientific (and metaphysical) problems about the nature and diversity of minds and how evolution was able to produce them.

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