NB: WORK IN PROGRESS
STILL CHANGING

Construction kits required for biological evolution

(Including evolution of minds and mathematical abilities.)

The scientific/metaphysical explanatory role of construction kits

(CHANGING DRAFT: Stored copies will soon be out of date.)

NB: this version was last modified on 19 March 2015.
On 18 Apr 2015, a new version replaced this here:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

Aaron Sloman
School of Computer Science, University of Birmingham.
(Philosopher in a Computer Science department)

Background note added, 18 Feb 2015,
then moved to separate document, 1 Mar 2015

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/keller-org.html
A few notes on Evelyn Fox Keller's papers on
Organisms, Machines, and Thunderstorms:
A History of Self-Organization

JUMP TO TABLE OF CONTENTS

Begun: 15 Dec 2014
(Based partly on earlier documents on the Meta-Morphogenesis project)
Last updated:
9 Mar 2015: continuing reorganisation. Not yet done.
7 Mar 2015: started more reorganisation. Not yet done.
Re-named mixed construction-kits "hybrid construction-kits" 6 Mar 2015
Major reorganisation and expansion: 18 Feb 2015
4 Feb 2015; 11 Feb 2015; 14 Feb 2015
1 Feb 2015: much reorganisation, gap-filling, pruning;
24 Jan 2015: Replaced Intermediate Construction Kit (ICK) with Derived Construction Kit (DCK); 28 Jan 2015; 31 Jan 2015
16 Jan 2015 (Reorganised. Abstract added.); 21 Jan 2015
8 Jan 2015 added link to new book by Andreas Wagner.
5 Jan 2015: copied to Slideshare.net; 6 Jan 2015
4 Jan 2015: reorganised, with many parts re-written and expanded, including speculations about the future of theoretical physics.

19 Dec 2014; 25 Dec 2014 (expanded); 1 Jan 2015

This paper is
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
From time to time a PDF version will be generated from the html, available as:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.pdf (May be out of date.)

A version of this document was posted to slideshare.net on 5 Jan 2015, making this available in flash format, updated occasionally. That version will not necessarily be updated whenever this one is. (Not all the links work in the slideshare version.) See
http://www.slideshare.net/asloman/construction-kits

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

NOTE: An earlier paper posted to this location on "Construction kits as explanations of possibilities", defending Chapter 2 of The Computer Revolution in Philosophy (1978) against criticisms by reviewers has now been moved to a different location:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html

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

---

CONTENTS

ABSTRACT
Modified 24 Jan 2015; 31 Jan 2015; 19 Feb 2015

This is part of the Turing-inspired Meta-Morphogenesis project, which aims to identify transitions in information-processing since the earliest proto-organisms, in order to provide new understanding of varieties of biological intelligence. Transitions depend on "construction-kits", including the initial "Fundamental Construction Kit" (FCK) based on physics and "Derived Construction Kits" (DCKs) produced by evolution, development, learning, and culture. Some construction kits (e.g. Lego, Meccano, plasticine, sand) are *concrete*: using physical components and relationships. Others (e.g. grammars, proof systems and programming languages) are *abstract*: producing abstract entities, e.g. sentences, proofs, and new abstract construction kits. Some are meta-construction kits: able to create, modify or combine construction kits. Construction kits are generative: they explain sets of possible construction processes, with mathematical properties and limitations. Evolution and development demonstrate new possibilities for construction kits: evolution as a "blind theorem prover", proving "theorems" about what is and is not possible for the kits used, including meta-cognitive abilities to think and reason about mathematical discoveries, and discuss them with others, leading, for example, to Euclid's *Elements*. FCKs and DCKs help to provide new answers to old philosophical questions, e.g. about the nature of mathematics, language, mind, science, and life, and expose deep connections between science and metaphysics. The requirement to show how the FCK makes *everything else* possible provides a challenge for physicists: demonstrate that your fundamental theory can explain how all the products of natural selection are possible. A core thread is the connection of control and (semantic) information. The aim is to explain, not reduce: even qualia are products of products of evolution -- and will occur in robots with animal intelligence. This paper merely introduces a large research programme that seems to have a chance of being progressive, in the sense of Lakatos.

INTRODUCTION
Background: What is science? Beyond Popper and Lakatos
(Background preamble)

NOTE:
Part of this introductory section is shared between two documents:
This document (on construction-kits):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
and a document on the role in science of explanations of possibilities:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html

I'll present a partial outline answer to the question: how was it possible for the known and unknown varieties of life to evolve from lifeless matter, including some varieties that are able to make the mathematical discoveries assembled in Euclid's *Elements*? The answer is based on construction kits, both fundamental and derived.

Karl Popper (1934) distinguished scientific and non-scientific statements. He required the former to be empirically falsifiable, otherwise they were metaphysical. Unfortunately this criterion has been blindly followed by many scientists who seem to be ignorant of the history of science. E.g. the ancient atomic theory of matter was not falsifiable, but was an early example of a deep scientific theory. Popper (unlike many who use falsifiability as a criterion for scientific content) acknowledged that some unfalsifiable metaphysical theories could be precursors of scientific theories, but it is arguable that labelling them "metaphysics" rather than "science" is misleading, because of their importance for science.
http://plato.stanford.edu/entries/democritus/#2
http://en.wikipedia.org/wiki/Democritus

Popper's philosophy of science was extended by Imre Lakatos (1980), who proposed ways of evaluating competing scientific research programmes, based on their progress over time. He distinguished "progressive" and "degenerating" research programmes, and specified criteria for deciding which of two research programmes is better, though it is not always possible to decide while both are being developed. The history of science shows that what appears to be a decisive victory (like Thomas Young's evidence of diffraction of light, which was taken to disprove Newton's particle theory of light) can later be overturned (e.g. when light was shown to have a dual wave-particle nature).

Chapter 2 of Sloman (1978) extended the ideas of Popper and Lakatos to accommodate scientific theories concerned with what is *possible*, e.g. types of plant, types of animal, types of reproduction, types of thinking, types of learning, types of verbal communication, types of molecule, types of chemical interaction, and types of biological information-processing -- the focus of the Meta-Morphogenesis project (Sloman (2012+)).

A separate "companion paper" discusses in more detail the general concept of 'explaining possibilities', its importance in science, the criteria for evaluating such explanations, and how this notion conflicts with the requirement for all scientific theories to be falsifiable. Further examples are in a closely related paper, also straddling science and metaphysics, on 'Actual Possibilities' (published in 1996) freely available online here.

Claiming that there are sharp boundaries between science and metaphysics harms both science and metaphysics. Both science and metaphysics can be pursued with rigour and openness to specific kinds of criticism. Some of the criteria for evaluating theories of what is possible, including theories that straddle science and metaphysics, were presented in sections 2.5.4 and 2.5.5 of Chapter 2 of the 1978 book, and in the section entitled Why allowing non-falsifiable theories doesn't make science soft and mushy in the companion paper to this one.

The ideas presented here are offered as a contribution to metaphysics as well as science -- but incomplete science, leaving much to be done. There are also deep (Kantian) connections with philosophy of mathematics, and explanations of necessity, that may seem surprising. Possible engineering applications include understanding the difficulties in designing machines with animal intelligence.

(End of preamble.)

JUMP TO TABLE OF CONTENTS

Fundamental and Derived Construction Kits (FCK, DCKs)

Life requires construction kits able to build machines with many capabilities, including growing highly functional bodies, immune systems, digestive systems, repair mechanisms, and reproductive machinery. The requirements for life include information processing (e.g. deciding what to repair) as well as physical construction. The Fundamental Construction Kit (FCK) provided by the physical universe when our planet came into existence was sufficient to make possible all the forms of life that have so far evolved on earth, including the challenges that helped to drive evolution of new forms of life. It also makes possible many unrealised but possible forms of life, in possible but unrealised types of physical environment. How does the FCK make all these things possible?

Figure **FCK**, below, indicates crudely how a common initial construction kit could explain many possible trajectories in which components of the kit are assembled to produce instances of possible living and non-living physical forms.

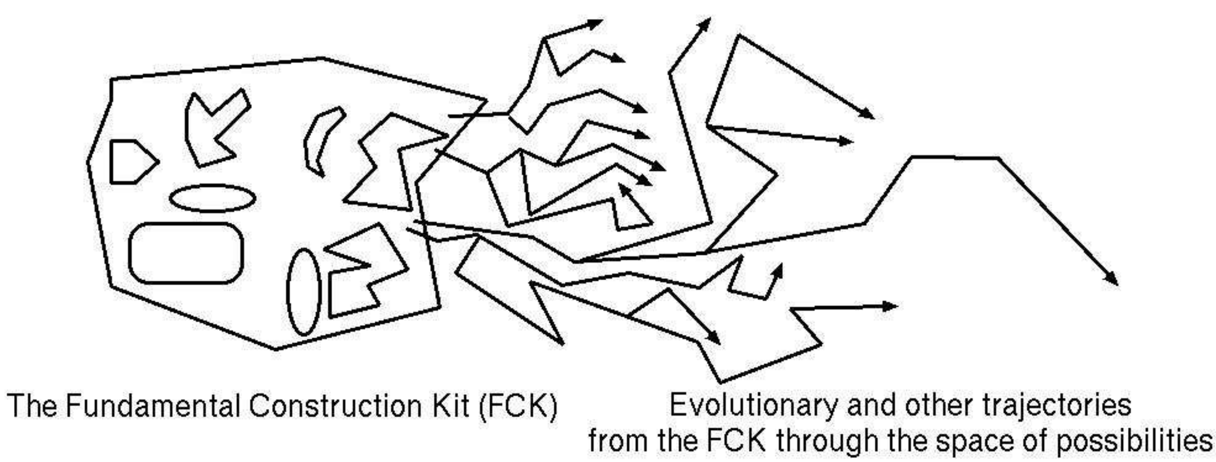

Figure FCK: Fundamental Construction Kit and possible trajectories
Think of time and increasing complexity going approximately from left to right.

Later, I'll return to discussion of how such a construction kit, and its products, can have mathematical features that make certain design abstractions useful for evolution, then later helped to drive evolution of mathematical competences in biological mechanisms, and eventually mathematical, and meta-mathematical, competences in some animals, a story that presents new biological/evolutionary foundations for mathematics.

The history of technology, science and engineering includes many transitions in which new construction kits were derived from old ones. That is true in particular in the science and technology of digital computation, where advances made use, for example, of

- punched cards, punched tape, and mechanical sorting devices,
- electronic circuits, switches, mercury delay lines, vacuum tubes, switchable magnets, and other devices,
- arrays of transistors, connected electronically,
- properties of machine language instructions expressed as bit-patterns, initially laboriously 'loaded' into a computer by setting banks of switches,
- symbolic machine languages composed of mnemonics that are "translated" by mechanical devices into bit-patterns on punched cards or tapes that can be read into a machine to get it set up to run a program,
- compilers and assemblers able to translate symbolic programs into required bit patterns,
- many types of higher level programming language that are compiled to machine language or intermediate level languages before programs start running,
- use of operating systems: programs that manage the running of other programs,
- higher level programming languages that are never *compiled* (i.e. translated into and replaced by programs in lower level languages) but are *interpreted* at run time, with each interpreted instruction triggering a collection of behaviours, possibly in a highly context sensitive way.

A vast number of related changes made it possible for human designers to produce larger, more complex, and more powerful systems, with the aid of increasingly complex tools for designing, building, testing, debugging, etc.

Products of evolutionary trajectories from the FCK may combine to form Derived Construction Kits (DCKs) (some specified in genomes, and some designed or discovered by individuals), that speed up construction of more complex entities with new types of properties and behaviours, crudely indicated in Figure DCK. In cases of convergent evolution, new DCKs with overlapping functionality, using different mechanisms, may evolve in different species in different locations. A DCK producing mechanisms enabling elephants to learn to use trunk, eyes, and brain to manipulate food may share features with a DCK enabling primates to acquire abilities to use hands, eyes, and brains to manipulate food. Both competences apparently evolved after the last common ancestor.

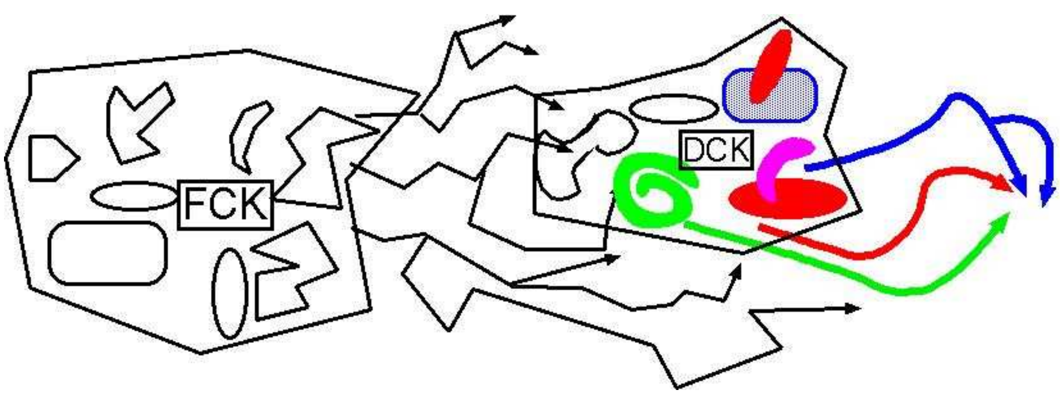

Figure DCK: Derived Construction Kit and new possible trajectories

The fundamental construction kit (FCK) on left can give rise to new evolved "derived" construction kits, such as the DCK on the right, from which new trajectories can begin, producing new more complex designs, e.g. organisms with new morphologies and new information processing mechanisms. The shapes and colours indicate (crudely) the presence of qualitative differences between components of the new construction kit and components of the original. Time again goes from left to right.

Biological evolution seems to have produced many branching lineages of increasingly complex re-usable construction kits, adding new, more complex, types of physical and chemical process (e.g. new forms of reproduction), and increasingly complex forms of information processing, among other products, some summarised below.

Details of human-designed forms of computation look very different from evolved biological layers of machinery for assembling complex information processing systems from simpler ones. But there may be deep similarities of function, including use of virtual machinery, discussed below. Over time, human designers use their evolved mechanisms, to produce larger, more complex, and more powerful systems, with the aid of increasingly complex tools for designing, building, testing, debugging, etc. Likewise evolution.

Some new biological construction kits allow creation of new physical materials with new properties -- e.g. different weight/strength ratios, different kinds of flexibility and elasticity, different sorts of permeability, different ways of storing, releasing and using energy, different ways of producing motion, different forms of reproduction, and many more, all making use of new chemical mechanisms, including products of "biological nano-engineering".

Different life-forms (microbes, fungi, slime moulds, plants of many sizes and shapes, invertebrate and vertebrate animals of many kinds) have produced different sorts of physical materials used in constructing bodies, or extensions of bodies such as webs, cocoons and egg-shells. Examples include the cellulose and lignin structures that play an important role in providing the strength of large plant structures that grow upwards out of soil, the materials in animals that produce rigid or semi-rigid structures (bones, shells, teeth, cartilage), the materials used in flexible structures with high tensile strength (e.g. tendons, vines), materials used in absorbing nutrients, oxygen, or water from the environment, materials transported between body parts, for different purposes (nutrients, waste matter, hormones, information, e.g. about stress or damage), materials concerned with storage and transfer or deployment of energy, for heat, for applying forces, for mobility, for reproduction, and many more.

Note on making possible
The assertion "X makes Y possible" does not imply that if X does not exist then Y is impossible. It may be possible for X to be built, making construction of Y easier. If Y is described at a suitable level of abstraction, then other things than X can make Y possible, for instance, an alternative construction kit. So "makes possible" should be interpreted in our discussion as a relation of sufficiency, not necessity. The exception is the case where X is the FCK -- the {\em Fundamental Construction Kit} -- since all concrete constructions must start from that. If X and Y are abstract, it is not clear that there is something like the FCK to which they must be traceable. The space of abstract construction kits may not have a fixed "root" kit. However, the abstract construction kits that can be thought about by physically implemented thinkers may be more constrained.

-- The variety of biological construction kits

As products of physical construction kits become more complex, with more ways of contributing to needs of organisms, and directly or indirectly to reproductive fitness, their use requires increasingly sophisticated control mechanisms, for which additional sorts of construction kit are required, including kits for building information-processing mechanisms of various sorts.

Compare (a) microbes that use only a few chemical sensors providing information about the *immediate* external and internal physical environment, with very few behavioural options, and (b) complex organisms that acquire and use information about enduring spatial locations in extended terrain with various resources (e.g. types of food) and dangers (e.g. noxious substances, lurking predators). The latter construct and use complex (internal or external) information stores about their environment, whereas the former merely acquire, use and replace fragments of information, using the same types of internal information throughout their life.

Another transition is to organisms that acquire and use information about information-processing, in themselves and in others, e.g. conspecifics, predators and prey. What sorts of construction kits suffice?

Some information processing systems are control systems made of components whose states are represented by physical measures, often referred to as "variables" and "constants", storing numbers that change discretely or continuously or not at all. Some represent states of sensors, others states of outputs, and others internal states of various sorts. In these control systems relationships between components are represented mathematically by equations, including differential equations, and possibly also constraints (e.g. inequalities) specifying restricted, possibly time-varying, ranges of values for the variables. Such a system with N variables has a state of a fixed dimension, N. The only way to store new information in such systems is in static or dynamic values for the variables -- changing "state vectors". Similar resources are required for modelling weather systems.

There are many well understood special cases of this pattern, such as simple forms of homeostatic control using negative feedback. Neural net controllers may be very much more complex with variables typically clustered into strongly interacting sub-groups, and perhaps groups of groups, etc.

Some fairly recent discoveries indicate that biological evolution made use of quantum-mechanical features of the FCK that are not yet fully understood, but suggest that there are forms of information processing that are very different from what current computers do. E.g. a presentation by Seth Lloyd, summarises roles of quantum phenomena used in deep sea photosynthesis, avian navigation, and odour classification.
https://www.youtube.com/watch?v=wcXSpXyZVuY

This may turn out to be the tip of an iceberg full of quantum-based information-processing mechanisms used in many DCKs not yet discovered by us, e.g. us (e.g. Hammeroff & Penrose (2014))

-- More varied mathematical structures

Partly as a result of use of computers, the variety of types of control in artefacts has exploded, using developments in logic, linguistics, and various parts of AI dealing with planners, learning systems, problem solving systems, vision systems, theorem provers, teaching systems, map-making explorers, automated circuit designers, program checkers, and many more. The world wide web can be thought of as an extreme case of a control system made up of millions of constantly changing simpler control systems, interacting in parallel with each other and with millions of display devices, sensors, mechanical controllers, humans, and many other things. The variety of types of control mechanism in computer-based systems extends far beyond the sorts familiar to control engineers, and studied in control theory. See
http://en.wikipedia.org/wiki/Control_theory
http://en.wikipedia.org/wiki/Nonlinear_control

Many different sorts of control system may be required in the life of a single organism, e.g. between an egg being fertilised and the death of the organism.

Not all natural control functions are numerical. A partially constructed percept, thought, question or plan for action, has parts and relationships, to which new components and relationships can be added, and others removed, as the construction proceeds and the product (percept, thought, plan, map) becomes more complex. There are different branches of numerical and non-numerical mathematics suited to the problem of designing or understanding such systems, including graph theory, lattice theory, knot theory, category theory, set theory, logic, and others.

For a full understanding of mechanisms and processes of evolution and development, new branches of mathematics are likely to be needed, including mathematics relevant to complex non-numerical changes in structural relationships, such as revising a grammar for internal records of complex structural information.

Traditional vector- and equation-based control theories (with probabilistic extensions) are not general enough for intelligent control systems that build and use sentences, problem descriptions, structural descriptions, explanatory theories, plans of varying complexity, learning mechanisms, systems of motives, values, social rules, and rule-based games, among other things. They almost certainly cannot adequately represent changes of molecular structure in which bonds are turned on or off Anderson (1972). It looks as if evolution, like human mathematicians and computer scientists millions of years later, built construction kits and information structures able to cope with structures of changing complexity, unlike the models and mechanisms based only on fixed sets of variables linked by equations -- without any means of representing either the structure of the meaning of a sentence, such as this one, or the structures of many perceived processes, including waves breaking on a rocky seashore or an intricately choreographed ballet.

It is unlikely that all the required forms of information, all the forms of control, and all the types of physical mechanism required for implementation are already understood by scientists and engineers. Yet the FCK along with the DCKs produced directly or indirectly by natural selection must be sufficiently general to model and explain everything that has evolved so far.

Clearly there is a huge variety of types of construction kit, that cannot all be surveyed here. In view of all the diversity and complexity of biological structures and processes, this paper includes no attempt at a complete theory. Its purpose is merely to present a research framework within which gaps in our understanding can be discovered and where possible filled, possibly over several decades, or even centuries. In particular, this is a first draft attempt at specifying some features of old and new construction kits, in the hope that later research will fill many of the gaps.

-- Non-uniform distribution of products of the FCK

I am not suggesting that the planet on its own could generate all those life forms. Life on earth depends crucially on energy from solar radiation (though future technologies may remove that dependence). Other external influences that were important for the particular forms of life that evolved on earth included asteroid impacts, and cosmic radiation. The role of entropy is discussed in: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html.

Before our solar system formed, the fundamental construction-kit was potentially available everywhere in the universe, making possible the formation of galaxies, stars, clouds of dust, planets, asteroids, and many other lifeless entities, as well as supporting all forms of life, possibly through derived construction kits (DCKs) that exist only in special conditions. However, local conditions e.g. extremely high pressures, temperatures, gravitational fields, distribution of kinds of matter, etc. may locally mask some parts of the FCK or prevent them from functioning.

According to some physical theories, every physical particle is (or can be) spread out over large areas, or possibly over the whole universe: nevertheless there must be differences in what exists in different places, for different processes can occur in different places.

So the contents of the FCK are not necessarily distributed uniformly throughout the universe and some developments based on the FCK are impossible in certain parts of the universe lacking the required matter, or other pre-requisites.

The FCK must in some sense be available at the centre of the sun, but that does not mean that animal life or plant life can exist there. Likewise if the cloud of dust from which the earth is thought to have formed had been composed mostly of grains of sand, then no DCK capable of supporting life as we know it could have emerged, since earth-life depends on the presence of carbon, oxygen, hydrogen, iron, and many other elements, whose distribution throughout the universe is not uniform.

As the earth formed, the new physical conditions created new DCKs that made the earliest life forms possible. A deep analysis of requirements for a DCK that supports primitive life forms is in Tibor Ganti's book The Principles of Life. The new DCK (building on the FCK) must have made possible both the formation of pre-biotic chemical structures and very simple life forms, and also the environments in which they could survive and reproduce. But there's more to life than primitive life forms!

There is a huge variety of types of construction kit, that cannot all be surveyed here. This work is still in its infancy and only very shallow discussions using a small number sub-cases can be offered here.

Construction kits that will not be discussed here but should be in a more complete investigation include internet-based virtual construction kits such as Minecraft (https://minecraft.net/) currently used by millions of people. Other sorts of virtual machinery will be mentioned later.

Construction kits generate possibilities and impossibilities
(Expanded 23/24 Feb 2015)

Explanations of how things are possible sometimes refer to construction kits, which may be manufactured, e.g. Meccano and Lego, or consist of naturally occurring materials, e.g. boulders, mud, or sand. (I am not suggesting that naturally occurring construction kits have clear boundaries, like manufactured kits.) Each kit makes possible certain types of construct, instances of which can be built by assembling parts provided in the kit. Some construction kits use products of products of biological evolution. For example, some birds' nests are assembled from twigs or leaves.

In some cases, properties of components, such as shape, are inherited by constructed objects. E.g. objects composed only of Lego bricks joined in the "standard" way all have external surfaces that are divisible into faces parallel to the surfaces of the first brick used. However, as Ron Chrisley pointed out to me, when two Lego bricks are joined at a corner only, using only one stud and socket, it is possible to have relative rotation, a possibility I had not noticed. (Similar attention failures can occur in mathematicians, as documented in Imre Lakatos (1980).)

More generally, constructed objects can have features none of the components have, e.g. a hinge is a non-rigid object that can be made from rigid objects: two rigid objects with aligned holes through which a rod or screw is passed, creating a flexible object from non-flexible parts. There are many such examples of emergent novelty. (I don't know if anyone has attempted an exhaustive taxonomy of ways of producing novel powers, structures and processes by combining old components in a new way.)

A construction kit that makes some things possible and others impossible can be extended so as to remove some of the impossibilities, e.g. by adding a hinge to Lego, or adding new parts from which hinges can be assembled. Another option is to recruit something outside the kit, e.g. a gravitational field. Something like a seesaw can be made using gravity (part of the FCK) discussed above) to keep one piece supporting another that behaves as if hinged at the centre.

Lego, meccano, twigs, mud, and stones, can all be used in construction kits whose constructs are physical objects occupying space and time, namely concrete construction kits. There are also abstract construction kits, whose products do not exist in space-time, for example components of languages, such as vocabulary and grammar, or methods of construction of arguments or proofs. Physical representations of such things, however, can occupy space and/or time, e.g. a spoken or written sentence, a diagram, or a proof presented on paper, or orally. There are also hybrid concrete+abstract construction kits, such as the physical board plus pieces of a chess set combined with abstract rules specifying legal moves, and conditions for winning and losing.

In some hybrid construction kits, such as chess, the physical pieces are not required by experts. For an expert, chess uses an abstract construction kit, and the physical components are dispensable for the purposes of playing chess, although communication of moves between players will need physical mechanisms. Moreover, the chess player's brain will use physical mechanisms (still not understood) to represent the abstract structures, states and processes. Closely related abstract structures, states and processes can also be implemented in computers, which can now play chess better than most humans. More cases are discussed below.

-- Construction kits for making information-users

Consider a construction kit consisting of a kind of material that can be deformed under pressure, e.g. plasticine or damp clay. If some object, e.g. a coin, or a rock, is pressed against a lump of the material the lump will change its shape, acquiring a new depressed portion whose surface has the inverted shape and size of part of the pressed object (e.g. a footprint). If two objects are pressed against different parts of the material then there will be two imprints.

If the two objects have the same shape, or the same maximum length, then the two imprints will also (if they are pressed in the same way). Should we therefore say that the clay not only retains information about the size and shape of things pressed into it, but also their similarity and differences, e.g. same shape but different lengths, or vice versa? Likewise if two small hands and two large feet are pressed into the clay should we say that the clay can sort the things pressed into it according to shape, and according to size?

The clay does not, in itself, have the ability to make use of those relationships, but if something else can inspect the clay it may be able to take decisions or answer questions about the things that were pressed into it, including quite abstract questions, e.g. about whether any two of the objects were similar in shape, or how they differ. This requires mechanisms able to make use of highly abstract concepts of sameness and difference. But we must be careful not to jump to conclusions from uses we can make of physical differences, as may happen when scientists discover changes in brain states associated with things for which we have labels.

One of the claims I am making is that evolution was able to assemble machines that could acquire and make use of various kinds of information about things in their immediate vicinity, then later extended those capabilities in dramatic ways. Making use, in the simplest cases, is selecting between alternative possible actions. A very simple case is a thermostat that turns a heater on or off. A more complex case is a rotary Watt governor that controls the amount of steam going to a rotary steam engine in a way that keeps the speed of the engine within a certain range.

There are many evolved mechanisms that acquire information through transducers of very different sorts and use the information in ways that produce, maintain or avoid various states of affairs. The role a physical state or change plays in controlling something, e.g. deployment of energy, or direction of growth, can be described as providing information, in this case *control* information. As Gibson (1966) pointed out, this often requires cooperation between processes of sensing and acting, as well as combined use of concrete and abstract construction kits.
Slime moulds are a spectacular example.
http://www.theguardian.com/cities/2014/feb/18/slime-mould-rail-road-transport-routes
The uses of motion in haptic and tactile sensing and the use of saccades, changing vergence, and other movements in visual perception are all examples of the interplay between sensing and doing, in "online intelligence". But there were many details Gibson seems to have ignored because they involve "offline intelligence", e.g. using perceptual information when reasoning about more or less remote possibilities.

The same information-bearing structure (e.g. the impression of a foot, the shape of a rock) can provide very different information to different information-users, depending at least on (a) what kinds of sensors they can use to get information from the structure, (b) what sorts of information-processing (storing, analysing, comparing, combining, synthesizing, retrieving, deriving, using...) mechanisms the users have, (c) what sorts of needs or goals they can serve (knowingly, or unwittingly) by using various sorts of information.

So, from the fact that changes in some portion of a brain are strongly correlated with changes in some aspect of the environment we cannot conclude much about what information about the environment the brain acquires and uses or how it does that -- any more than discovering footprints in the sand where animals walk, tells us that a beach perceives animals.

During evolution, and to some extent during individual development, all of (a), (b) and (c) become more complex, more diversified, and more capable of going on in parallel while they interact (sometimes competing, sometimes collaborating, sometimes invoking, extending, recording, controlling, redirecting, enriching, training, abstracting, refuting, or terminating) within a complex information processing architecture. If we don’t understand the architecture and the many information-processing functions it supports, and how they are related, we’ll understand very little about the biological functions of observed fragments. That architectural knowledge will not be expressible in sets of differential equations, or statistical relationships. (For impressive but partial attempts to characterise some architectural roles in human information processing see Minsky (1987), Minsky (2006). Compare Sloman (2003).)

There is much we still do not know about the construction kits used, and what they are used for. The aim of the Meta-Morphogenesis project is to investigate the huge variety of cases of use of information, partly because it is possible that there are some uses that we have not noticed, which are essential for understanding the more complex control phenomena in living things, including understanding more of the things brains do. It is very likely that the assumptions currently made by neuroscientists about the information processing in brains omit some of the most important types, and that AI researchers influenced by those assumptions are therefore failing to replicate some important functions of brains in their robots and other machines.

It is also very likely that progress in this investigation will require major conceptual advances regarding what the problems are and what sorts of answers are relevant. E.g. "Where in the brain are discoveries made?" "Where do emotions occur in the brain?" "Where in the brain is musical ability?"

"Where does understanding occur when you read a sentence?" may all turn out to be nonsensical questions. But that does not mean there are no mental states and processes, including detection of changes in qualia.

-- Biological construction kits

How did the FCK support more complex life forms? Is the Darwin-Wallace theory of natural selection the whole answer? Graham Bell wrote in Bell (2008)

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

No: an adequate explanation must mention both **selection** mechanisms and **generative** mechanisms. Without generative mechanisms, selection processes will not have a supply of new viable choices, and the selection mechanisms will not be able to select new options. Moreover, insofar as environments providing opportunities, challenges and threats are part of the selection process, the construction kit used by evolution requires mechanisms not intrinsically concerned with life, e.g. volcanoes, earthquakes, asteroid impacts, lunar and solar tides, and many more, in addition to evolved construction kits and their products.

The idea of evolution producing construction kits is not new, though they are often referred to as "toolkits". Coates et al. (2014) ask whether there is "a genetic toolkit for multicellularity" used by complex life-forms. Toolkits and construction kits normally have users (e.g. humans or other animals), whereas the construction kits we have been discussing (FCKs and DCKs) do not all need separate users.

Both selection mechanisms and generative mechanisms change during evolution (partly by influencing each other). Natural selection (blindly) uses the initial enabling mechanisms provided by physics and chemistry not only to produce new organisms, but also to produce new richer DCKs, including increasingly complex information-processing mechanisms. Since the mid 1900s, spectacular changes of this sort have happened in human-designed computing mechanisms, including new forms of hardware and new forms of virtual machinery, and networked social systems all of which would be incomprehensible to early hardware designers. Similar changes during evolution produced new biological construction kits whose products are even less comprehensible to scientists (and philosophers) familiar only with the initial mechanisms and what they can observe.

Biological DCKs have made possible not only a huge variety of physical forms, and physical behaviours, but also forms of *information processing* required for increasingly complex control problems, as organisms become more physically complex and more intelligent in coping with their environments, including interacting with predators, prey, mates, offspring, conspecifics, etc. In humans, that includes abilities to form scientific theories and discover and prove theorems in topology and geometry, some of which are used unwittingly in practical activities. I suspect many animals come close to this in their systematic (but unconscious) abilities to perform complex actions in a way that makes use of mathematical features of environments. Nest building abilities and abilities involved in hunting and consuming prey may overlap with the topological and geometrical competences of humans. (See the discussion of below).

Concrete (physical), abstract and hybrid construction kits

Products of a construction kit may be concrete, i.e. physical, or abstract, like a proof, a sentence, or a symphony; or hybrid, e.g. a physical presentation of a proof or poem.

Concrete kits:
Construction kits for children include physical parts that can be combined in various ways to produce new physical objects that are not only larger than the initial components but have new shapes and new behaviours. Those are *concrete* construction kits. The FCK is a concrete construction kit, whose details are the subject of much research by physicists.

Abstract kits:
Despite the current (deeply confused) fashion emphasising embodied cognition, many examples of thinking, perceiving, reasoning and planning, require abstract construction kits. For example, planning a journey to a conference does not require physically trying possible actions, in anything like the way water finding a route to the sea explores possible route-fragments. Instead an abstract construction kit representing possible options and ways of combining them can be used. Being able to talk requires use of a grammar that specifies a collection of abstract structures that can be assembled using a collection of grammatical relationships to form new abstract structures with new properties relevant to various tasks involving use of information. The sentences allowed by a grammar for English can be thought of as abstract objects that can be instantiated in written text, printed text, spoken sounds, sign languages, morse code, and other concrete forms. In that sense a grammar is an abstract construction kit whose constructs can have concrete (physical) instances.

A grammar does not specify a language: a *semantic* construction kit, structurally related to the grammar, is required for building possible *meanings* for the language to express. The use of a language depends on the existence of language users, for which more complex construction kits are required, also products of evolution and learning. (Evolution of various types of language is discussed in this presentation, which argues that internal languages must have evolved first, then sign languages.)

Hybrid abstract+concrete kits:
These are combinations, e.g. physical chess board and chess pieces combined with the rules of chess, lines and circular arcs on a physical surface instantiating Euclidean geometry, puzzles like the mutilated chess-board puzzle, and many more. A particularly interesting hybrid case is the use of physical cubes to instantiate arithmetic, which may lead to the discovery of prime numbers when certain attempts at rearrangement fail -- and an explanation is found.

In computing technology, physical computers, programming languages, operating systems and virtual machines can be thought of as hybrid construction kits that can make things happen when they run Sloman (2013). A logical system with axioms and inference rules can be thought of as an abstract kit supporting construction of logical proof-sequences combined with a physical notation for written or printed proofs. A logical system cannot have physical causal powers whereas its concrete instances can, e.g. helping a student separate proofs and non-proofs.

Natural selection seems to have "discovered" the power of hybrid construction kits, especially the use of sophisticated virtual machinery, long before human engineers did, though probably in much more powerful forms than current engineering designs Sloman (2010). All examples of perception, learning, reasoning, and intelligent behaviour include hybrid construction kits, though scientific study of such kits is still in its infancy. This discussion merely scratches the surface of a huge multi-disciplinary research area. Work done so far on the Meta-Morphogenesis project suggests that natural selection "discovered" and used a staggering variety of types of hybrid construction kits that were essential for reproduction, for developmental processes (including physical development and learning), for performing complex behaviours, and for social/cultural phenomena. Jablonka and Lamb (2005) seem to come close to making this point, though they use different terminology.

Kits including external sensors and motors

Some toys interact with the environment by moving parts, e.g. wheels. A simple toy car may include a spring that can be wound up. When started the potential energy in the spring is transformed into mechanical energy via gears, axles and wheels that are in contact with external surfaces. Further interactions, altering the direction of motion, may result from collisions with fixed or mobile objects in the environment.

Some construction kits allow assembly of such toys. More sophisticated kits include sensors that can be used to provide information for an internal mechanism that uses the information to take decisions concerning deployment of available energy, for instance using light, sonar, or in the case of rats, using whiskers, to gain information that allows frequent changes of direction or speed of motion, e.g. in order to avoid collisions, or in order to move towards a source of electrical or chemical energy when internals supplies are running low. Some examples are provided in Braitenberg (1984), though he (or at least some of his admirers) unfortunately over-interpreted his vehicles as being capable of love, fear, etc. The ideas are demonstrated here:
http://www.it.bton.ac.uk/Research/CIG/Believable%20Agents/

In some cases the distinction between internal and external components is arbitrary. For example, a musical box may perform a tune under the control of a rotating disc with holes or spikes that cause a tone to be produced when they reach a certain location, during the rotation. The disc can be thought of as part of the music box. It can also be thought of as part of a changing environment, in which case the devices that detect the holes or spikes are external sensors.

If a toy train set has rails or tracks used to guide the motion of the train as it moves, then the wheels of the train can be thought of as sensing the environment and causing changes of direction in the train. This is partly like and partly unlike a toy vehicle that uses an optical sensor linked to a steering mechanism, so that a vehicle can follow a line painted on a surface. The railway track provides both the information about where to go and the forces required to change direction. The painted line, however, provides only the information, and other parts of the vehicle have to supply the energy to change direction, e.g. an internal battery that powers sensors and motors. Evolution uses both sorts: e.g. wind blowing seeds away from parent plants and a hunter following a scent trail left by its prey. An unseen wall uses force to stop your forward motion in a dark room, whereas a perceived wall provides information, not force, that causes you to decelerate (Sloman, 2011).

-- Mechanisms for storing, transforming and using information

Some information is acquired, used, then lost because it is immediately over-written, e.g. sensor information in simple servo-control systems with "online intelligence", where only the latest sensed state is used for deciding whether to speed something up, slow it down, change direction, start to grasp, etc. In more complex control systems, with "offline intelligence", some sensor information is saved, possibly combined with other previously stored information, and remains available for use on different occasions for different purposes. In the second case, the underlying construction-kit needs to be able to support stores of information that grow with time and can be used for different purposes at different times. Sometimes a control decision at one time can use items of information obtained at several different times and places, for example information about properties of a material, where it can be found, and how to transport it to where it is needed. Sensors used online may become faulty or require adjustment. Evolution may provide mechanisms for testing and adjusting. When used offline, stored information may need to be checked for falsity caused by the environment changing, as opposed to sensor faults.

There are hugely varied ways of acquiring and using information, some of which have been discovered (or re-discovered) and modelled by AI researchers, psychologists, neuroscientists, biologists and others, though it seems that evolution has achieved a great deal more, not only in humans, but in other intelligent animals. Many of these achievements require not just additional storage space but very different sorts of information-processing architectures. A range of possible architectures is discussed in Sloman (1993), Sloman (2003), and Sloman (2006). Some types use sub-architectures that evolved at different times, meeting different needs, in different biological niches Sloman (2000).

This raises the question whether evolution produced "architecture kits" able to combine evolved information-processing mechanisms in different ways, long before software engineers discovered the need. Such a kit could be particularly important for individuals that produce new sub-systems, or modify old ones, during individual development, e.g. during different phases of learning by apes, elephants, and humans.
The BICA society aims to bring together researchers on biologically inspired cognitive architectures. Some examples are here: http://bicasociety.org/cogarch/

-- Mechanisms for controlling position, motion and timing

All of the concrete construction kits (and some of the hybrid kits) share a deep common feature insofar as their components, their constructs and their construction processes involve space and time, both during construction processes, as items are moved together and their relationships altered, and during the behaviour of complex constructed objects. Those behaviours include both relative motion of parts of an object, e.g. wheels rotating, joints changing angles, and also motion of the whole object relative to other objects, e.g. an ape grasping a berry.

A consequence of the common spatiality is that objects built from different construction kits can interact, by changing their spatial relationships (e.g. if one object enters, encircles or grasps another), by applying forces that are transmitted through space, and in other ways. Products of different kits can interact in more complex ways, e.g. one being used to manipulate another, or one providing energy or information for the other.

Moreover, new construction kits can be formed by combining two or more concrete kits. In some cases this will require modification of a kit, e.g. supporting combinations of lego and meccano by adding pieces with lego studs or holes alongside meccano sized screw holes.

Another consequence of the fact that objects exist in space/time is the need for timing mechanisms. Organisms use many "biological clocks" operating on different time-scales controlling repetitive processes, including daily cycles, heart-beats, breathing, and wing or limb movements required for locomotion. More subtly there are adjustable speeds of motion or change, and adjustable rates of change. Examples: a bird in flight approaching a twig on which it is to land; an animal running towards a tree to escape a predator and having to decelerate as it approaches the tree to avoid a dangerous crash; a hand moving to grasp a stationary or moving object, with motion controlled by varying coordinated changes of joint angles at waist, shoulder, elbow and finger joints so as to bring the grasping points on the hand into a suitable location relative to the selected grasping points on the object. (The last example is still very difficult for robots, when grasping novel objects in novel situations: partly because of designs specifying the wrong ontologies.)

There are also mechanisms for controlling or varying rates of production of chemicals (e.g. hormones). So biological construction kits need many mechanisms with abilities to measure time intervals and to control rates of repetition or rates of change of parts of the organism. These construction kits may be combined with other sorts of construction kits that require temporal as well as spatial control, e.g. changing speed and direction of motion simultaneously. There are different requirements for controlling growth of fixed structures, e.g. trees growing branches, and for mobile animals.

-- Combining construction kits

At the molecular level there is now a vast, and rapidly growing, amount of research on interacting construction kits, for example interactions between different parts of the reproductive mechanism during development of a fertilised egg, interactions between invasive viral or bacterial structures and a host organism, and interactions with chemicals produced in medical research laboratories, among many other types.

In the realm of digital computation the ways of combining different toolkits include the application of functions to arguments, although both functions and their arguments can be far more complex than the simple cases most people encounter in learning about arithmetic. For example a function could be a compiler, its arguments could be arbitrarily complex programs in a high level programming language, and the output of the function in each case might be either a report on syntactic errors in the input program, or, if there are no errors, a machine code program to run on a particular type of computer.

The application of functions to arguments is a very different process from assembling structures in space time. If computers are connected via digital to analog interfaces, linking them to other things, e.g. surrounding matter, or if they are mounted on machines that allow them to move around in space and interact, that adds a kind of richness that goes beyond application of functions to arguments.

That additional richness is present in the modes of interaction of chemical structures which include both digital (on/off chemical bonds) and continuous changes in relationships, as discussed by Turing in his paper on the chemical basis of morphogenesis Turing (1952) (the paper that inspired the Meta-Morphogenesis project.)

-- Combining abstract construction kits

The possibility of combining concrete construction kits results from the fact that their instances occupy space and time. Combining abstract construction kits is not so straightforward. A simple example is combining letters and numbers to form coordinates for squares on a chess board, e.g. "a2", "c5", etc. More complex examples include combining notations for a human language and a musical system for writing songs, or combining a computer operating system (e.g. Linux) with a programming language (e.g. Lisp).

In living organisms, there are interactions between products of the same or different kits that involve information, e.g. use of information for sensing, predicting, explaining or controlling, including information about information (Sloman, 2011)

Researchers on systems combining many kinds of functionality have found it useful to design information-processing architectures that provide frameworks for combining different mechanisms and information stores. This is particularly important in large projects where different research groups are working on sensors, learning mechanisms, motor subsystems, reasoning systems, motivational systems, various kinds of meta-cognition, etc., with associated sets of tools supporting processes of design, implementation, testing, debugging. Our own SimAgent toolkit is one among very many:

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

Some of the common principles include the need to be able to support different sorts of virtual machines with causal interactions between them, and the physical environment, as explained in this tutorial overview: Sloman (2013).

It seems that in addition to discovering design patterns for physical mechanisms biological evolution also discovered re-usable frameworks for assembling complex *information processing* architectures, accommodating multiple interacting virtual machines, with modifications developed by different species -- including humans Minsky (1987), Minsky (2006). This is a topic for further research, which I have argued elsewhere will provide new insights into complex mental states and processes, including forms of self-consciousness, varieties of affective states, and processes of cognitive development, helping to explain processes of mathematical development.
Discussed in connection with "toddler theorems" in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html
illustrated briefly in connection with the process of putting on a shirt, here. (Contributions from observant parents and child-minders are welcome. Deep insights come from individual developmental trajectories rather than statistical patterns of development across individuals.)

Adding a new DCK can make some of the possible further developments quicker to reach -- fewer additional steps are required than were originally required, and the total search space for a suitable sequence of steps to a solution may be considerably reduced. This is partly analogous to the role of previously proved theorems in a new proof. Using previous results can considerably shorten a proof, and have a dramatic effect on the size of the search-space when searching for a proof. If the number of steps to a solution has been reduced by 10 and there are two options at every step, the search for a complete design may have been reduced by a factor of 2^{10}, i.e. 1024: reducing the remaining evolutionary search space required by a factor over a thousandfold -- if there is a solution in the remaining search space. Evolutionary search spaces are very much larger, and in principle design re-use could have an even larger impact on search spaces. Similarly, the ability to re-use modified versions of useful designs could dramatically reduce an evolutionary search space -- if there is a solution in the remaining search space.

Creating a new construction kit requires use of an appropriate meta-construction kit able to extend or modify or combine previously created construction kits. Evolution needs to be able to create new meta-construction kits using natural selection. The great creator/meta-creator is now spectacularly aided and abetted by its products, especially humans and their products!

Construction kits built during individual development (Genetically meta-configured, not pre-configured)

Some new construction kits are products of the process of evolution of a species and are shared between all members of the species (barring genetic abnormalities), alongside construction kits shared between species, such as those used in mechanisms of reproduction and growth in related species. But evolution has also discovered the benefits of what might be called "meta-construction-kits", namely mechanisms provided for members of a species that allow individuals to build new construction kits during their own development.

Examples include mechanisms for learning that are developed by individuals on the basis of their own previously encountered learning experiences, which may be different in different environments for members of the same species. Human language learning is a striking example: things learnt at earlier stages make new things learnable that might not be learnable by an individual transferred from a different environment, having experienced a different language.

This contrast between genetically specified and individually built capabilities for learning and development was labelled a difference between "pre-configured" and "meta-configured" competences in Chappell and Sloman (2007), summarised below. Mathematical development in humans seems to be a special case of growth of meta-configured competences.

The construction kits used for assembly of new organisms that start as a seed or an egg enable many different processes in which components are assembled in parallel, using abilities of the different sub-processes to constrain one another. As far as I can tell, nobody knows the full variety of ways in which parallel construction processes can exercise mutual control in developing organisms. One implication is that there are not simple correlations between genes and organism features.

Turing’s (1952) examples of diffusing chemicals causing patterns when they interact include only formation of superficial 2-D patterns. Explaining the different ways in which features of a genome can directly or indirectly orchestrate many parallel processes of growth, development, formation of connections, etc. is a difficult challenge. A possible framework for allowing abstract specifications in the genome to interact with details of the environment in instantiating complex designs is illustrated schematically in Figure EVO-DEVO below. This generalises Waddington’s "epigenetic landscape" metaphor in Waddington (1957), by allowing individual members of a species to partially construct their own epigenetic landscapes instead of merely following paths in a landscape that is common to the species. Related ideas are in Karmiloff-Smith (1992). Some of the implications of these ideas for attempts to understand genetic abnormalities such as autism are discussed in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/autism.html

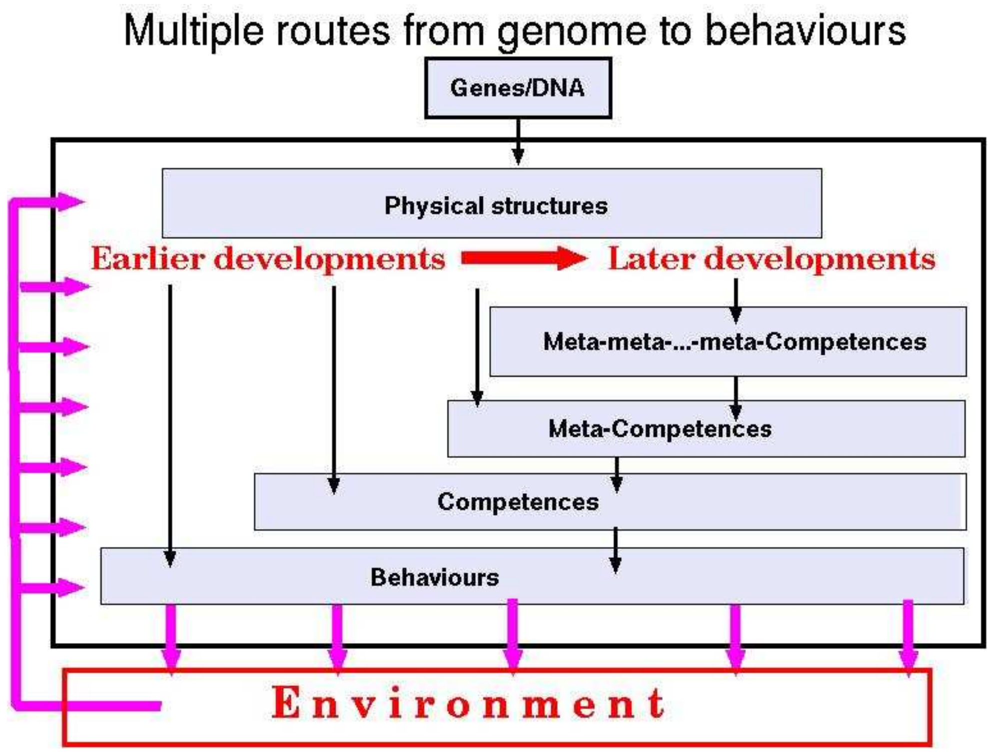

Figure EVO-DEVO: (Added: 22 Jan 2015)
Construction kits provided in individual genomes may give rise to very different individuals if the genome interacts with the environment in increasingly complex ways during development. Precocial species use only the routes on the left. Members of altricial species, using staggered development, may

show much greater variation in competences. Preconfigured competences are produced by the downward arrows on the left. Later, results of use of those competences interact with the genome to produce meta-configured competences on the right. (From Chappell and Sloman 2007)

Some constructions exclude or necessitate others

Physical construction kits (e.g. Lego, Meccano, Tinkertoys, plasticine, or a combination of paper, scissors and paste) have parts with physical properties (e.g. rigidity, strength, flexibility, elasticity, adhesion, etc.), possible relationships between parts and possible processes that can occur when the parts are in those relationships (e.g. rotation, bending, elastic or inelastic resistance to bending).

Features of a physical construction kit -- including the shapes and materials of the basic components, the ways in which the parts can be assembled into larger wholes, the kinds of relationships between parts and the processes that can occur involving them -- all contribute to explaining the possibility of entities that can be constructed from those components, and the possibility of processes, including both the processes of construction and the behaviours of the constructs.

Construction kits can also explain necessity and impossibility. A construction kit that has a very large set of generative powers initially can be used to build a structure realising some of the kit's possibilities, in which some further possibilities are excluded, namely all configurations that do not include what has so far been constructed. Some of the extensions that were possible before the last addition become impossible unless the last step is undone.

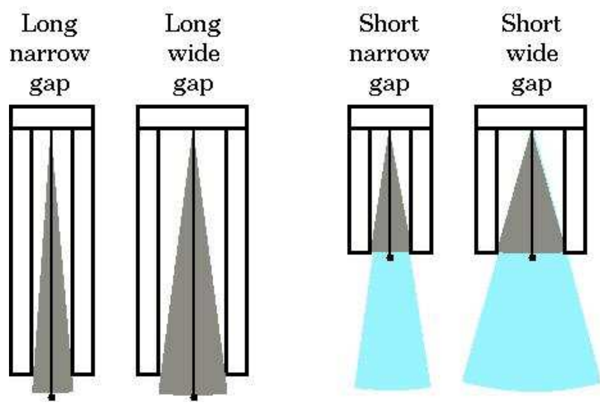

Figure GAPS:
Interactions between structure and remaining possibilities:

If a rod that can swing about a point in a plane is in a gap, then the wider the gap the wider the possible swing, and the shorter the rod for a fixed size gap, the wider the possible swing.

Moreover, what has been done may make some further steps possible and others impossible: e.g. the size of a gap between two rigidly assembled components will make it impossible to extend the structure by placing some components in the gap: A beam of 20cm square cross section cannot fit in a 10cm gap. Narrower beams can fit in the gap, but the angles by which their orientations can vary will depend on their diameter, the diameter of the gap, and other spatial relations. the narrower or shorter a beam in the gap, is the wider the angle through which it can rotate in a plane through the gap. The wider the gap is the wider the angle through which a beam of a certain width can rotate, while the longer the gap is the narrower the angle of rotation possible in that plane. Examples are in Figure GAPS. Both human engineers and evolution can make use of similar, though usually more complex, mathematical relationships, in skeletal geometry for example.

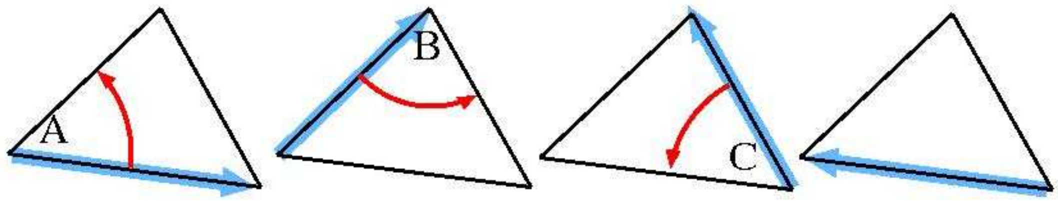

Figure TRIANGLE:
Mary Pardoe's proof of the triangle sum theorem.

The sequence of figures, demonstrates how the three-cornered shape has the consequence that summing the three angles necessarily produces half a rotation (180 degrees). Since the position, size, orientation, and precise shape of the triangle can be varied without affecting the possibility of constructing the sequence, this is a proof that generalises to any planar triangle. This is an unpublished proof reported to me by Mary Pardoe in the early 1970s. For more on this see
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html
Figure Triangle illustrates a different sort of example, where no physical properties of a structure (e.g. rigidity or impenetrability of materials) are involved, only spatial relationships. It presents a proof, found by Mary Pardoe, that internal angles of a triangle sum to a straight line, or 180 degrees. Unlike the "standard" proofs this proof makes no reference to Euclid's parallel axiom. The human mathematical ability to look at a physical situation, or a diagram representing a class of physical situations, and reason about constraints on a class of possibilities sharing certain constraints may have evolved from earlier abilities to reason about changing affordances in the environment Gibson (1979). Current AI perceptual and reasoning systems do not yet have these abilities, though that may change.

These are simple examples of the mathematical properties of construction kits (partly analogous to mathematical properties of formal deductive systems and AI problem solving systems). As parts (or instances of parts) of the FCK are combined, structural relations between components of the kit have two opposed sorts of consequences: they make some further structures possible, and they make other structures impossible -- and their absence or opposites, e.g. geometrical or topological properties, will then be necessary consequences of previous selection steps.

These examples illustrate how a construction kit with mathematical relationships can provide the basis for necessary truths and necessary falsehoods in some constructions (as in Sloman (1962), chapter 7). Such relationships between possibilities provide a deeper, more natural, basis for understanding modality (necessity, possibility, impossibility) than so called "possible world semantics". Since our examples of making things possible or impossible, or changing ranges of possibilities are examples of causation, this also provides the basis for a Kantian notion of causation based on mathematical necessity (Kant 1781), so that not all uses of the notion of "cause" are Humean (i.e. based on correlations), even if some are. Compare section Mathematical Discoveries, below.
For more on Kantian vs Humean causation see the presentations on different sorts of causal reasoning in humans and other animals, by Chappell and Sloman at the Workshop on Natural and Artificial Cognition (WONAC, Oxford, 2007):
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac

Varieties of causation that do not involve mathematical necessity, only probabilities (Hume?) or propensities (Popper) will not be discussed here.

-- Proof-like features of evolution

An unknown subset of the FCK, or perhaps a DCK or collection of DCKs, produced fortuitously as a side effect of formation of the earth, supported (a) primitive life forms and (b) processes of evolution that produced more and more complex forms of life, including new, more complex, derived, DCKs. New products of natural selection can make more complex products more reachable, as with toy construction kits. Moreover, there was not just one sequence of DCKs: different evolutionary lineages that evolved in parallel produced different DCKs. According to the "Symbiogenesis" theory, different DCKs produced independently can sometimes merge to support new forms of life combining different evolutionary strands. See http://en.wikipedia.org/wiki/Symbiogenesis

So creation of new DCKs in parallel evolutionary streams with combinable products can hugely reduce part of the search space for complex designs, at the cost of excluding parts of the search space reachable from the FCK. For example, use of DCKs in the human genome may speed up development of language and typical human cognitive competences, while excluding the possibility of "evolving back" to microbe forms that might be the only survivors after a cataclysm.

-- Euclid’s construction kit

A much older example, of great significance for philosophy of mathematics, is the construction kit specified in Euclidean geometry, starting with points, lines, surfaces, and volumes, and methods of constructing new more complex geometrical configurations using a straight edge for drawing straight lines in a plane surface, and a pair of compasses, for drawing circular arcs in a surface.

A different sort of geometry allows line segments to be translated and rotated in a plane while preserving their length. This is an assumption underlying the use of rulers for measuring length. Adding movable and rotatable line segments to Euclidean geometry allows an arbitrary angle to be divided into three equal parts, which is not possible in standard Euclidean geometry. See http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html A related construction is possible using "Origami geometry".

The idea of spaces of possibilities generated by different sorts of physical construction kits may be easier for most people to understand than the comparison with generative powers of grammars or formal systems, though the two are closely connected, since grammars and axiom systems are both abstract construction kits that can be parts of hybrid construction kits.

Concrete construction kits corresponding to them can be built out of physical structures: for example a collection of small squares with letters and punctuation marks can be used to form sequences that correspond to the words in a lexicon. Adding some blank squares and specifying rules of a grammar based on that lexicon can then be applied to sequences of squares, with blanks as word-separators, generating a set of possible physical sentences conforming to the grammar. The use of cursive ("joined up") script provides a more complex physical construction kit.

Some challenges for construction kits used by evolution, and also challenges for artificial intelligence and philosophy, arise from the need to explain both how natural selection makes use of mathematical properties of construction kits related to geometry and topology, in producing organisms with spatial structures and spatial competences, and also how various subsets of those organisms developed specific topological and geometrical reasoning abilities used in controlling actions and solving problems, and finally how at least one species developed abilities to reflect on the nature of those competences and eventually, through unknown processes of individual development and social interaction, using unknown representational and reasoning mechanisms, managed to produce the rich,

deep and highly organised body of knowledge published as Euclid's *Elements*. There are important aspects of those mathematical competences that as far as I know have not yet been replicated in Artificial Intelligence or Robotics, some of them listed in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/mathstuff.html
Is it possible that currently understood forms of digital computation are inadequate for the task, whereas chemistry-based information processing systems used in brains are richer?

-- Mathematical discoveries based on exploring construction kits

Some mathematical discoveries result from observation of naturally occurring physical construction kits and noticing how constraints on modes of composition of components generate constraints on resulting constructs. E.g. straight line segments on a surface can be joined end to end to enclose a region of the surface, but that is impossible with only two lines, as noted in (Kant, 1781). Likewise flat surfaces can be combined to enclose a volume, such as a tetrahedron or cube, but it is impossible for only three flat surfaces to enclose a finite space. It is not clear how humans detect such impossibilities: no amount of trying and failing can establish impossibility.

Many mathematical domains (perhaps all of them) can be thought of as sets of possibilities generated by construction kits of various kinds. Engineers deal with hybrid concrete and abstract construction kits. The space of possible construction kits is also an example, though as far as I know this is not a domain that has been explored systematically by mathematicians, though many special cases have.

In order to understand how the sorts of biological evolution that occurred on this planet are possible we need to understand the sorts of construction kits made possible by the existence of the physical universe, and in particular the variety of construction kits inherent in the physics and chemistry of the materials of which our planet was formed, along with the influences of its environment (e.g. solar radiation, asteroid impacts). An interesting research question is whether every construction kit capable of producing all the non-living structures on this planet would also suffice for evolution of all the forms of life on the planet, or whether life and evolution have additional requirements, e.g. external influences such as cosmic radiation.

Insofar as construction kits have mathematical properties, life and mathematics are closely interconnected, as we have already seen. More complex relationships arise after evolution of mathematical meta-cognitive mechanisms.

-- Evolution's (blind) mathematical discoveries

On the way to achieving those results, natural selection often works as "a blind theorem-prover". The theorems are mainly about new possible structures, processes, organisms, ecosystems, etc. The proofs that they are possible are implicit in the evolutionary trajectories that lead to such occurrences.

Proofs are often thought of as abstract entities that can be represented physically in different ways (e.g. using different formalisms) for the purpose of communication or persuasion (including self-persuasion), predicting, explaining and planning. It can also be argued that a physical sequence produced unintentionally, e.g. by natural selection, or by growth in a plant, that leads to a new sort of entity is a sort of (unwitting) proof that some construction kit makes that sort of entity possible. The evolutionary or developmental trail answers the question: how is that sort of thing possible? In that sense biological evolution can be construed as a "blind theorem prover", despite there being no intention behind the proof. Proofs of impossibility (or necessity) raise more complex issues, to be discussed elsewhere.

These observations seem to support a new kind of "Biological-evolutionary" foundation for mathematics, that is closely related to Immanuel Kant's philosophy of mathematics in his *Critique of Pure Reason* (1781), and my attempt to defend his ideas in my 1962 DPhil thesis (*Knowing and Understanding*, Oxford). This answers questions like "How is it possible for things that make mathematical discoveries to exist?", an example of explaining a possibility.

Those who try to go too directly from hypothesized properties of the primordial construction kit to explaining advanced capabilities such as human self-awareness (e.g. Schroedinger, Penrose) are likely to fail, because short-cuts will omit essential details of both the problems and the solutions, like mathematical proofs with gaps.

The success of many of the "mathematical discoveries" (or inventions?) produced (blindly) by evolution, depend on mathematical properties of physical structures or processes or problem types, whether they are specific solutions to particular problems (e.g. use of negative feedback control loops), or new construction-kit components that are usable across a very wide range of different species (e.g. the use of a powerful "genetic code", the use of various kinds of learning from experience, the use of new forms of representation for information, use of new physical morphologies to support sensing, or locomotion, or consumption of nutrients etc.)

These mathematical "discoveries" (discussed in more detail in other parts of the Meta-Morphogenesis web site) started happening long before there were any humans doing mathematics (which refutes Wittgenstein's suggestion that mathematics is an anthropological phenomenon). Many of the discoveries were concerned with what is possible, either absolutely or under certain conditions, or for a particular sort of construction-kit.

Other discoveries, closer to what are conventionally thought of as mathematical discoveries, are concerned with limitations on what is possible i.e. necessary truths.

Some discoveries are concerned with probabilities derived from statistical learning, but I think the relative importance of statistical learning in biology has been vastly over-rated because of misinterpretations of evidence. (To be discussed elsewhere.) In particular the important discovery that something is possible does not require collection of statistics: A single instance suffices. And no amount of statistical evidence can show that something is impossible.

For human evolution, a particularly important subclass of mathematical discoveries has been unwitting discovery and use of mathematical structures in the environment, a discovery process that starts before children are aware of what they are doing, and in some cases even before uses of language for communication have developed. Examples are discussed in the "Toddler Theorems" document (toddler theorems).

JUMP TO TABLE OF CONTENTS

Varieties of Derived Construction Kit

Evolution and its products use the fundamental construction kit of physics and chemistry to produce derived construction kits with different properties, including concrete, abstract and hybrid construction kits. These DCKs differ (a) at different stages within a lineage, (b) across lineages (e.g. in different coexisting organisms such as plants, insects, vertebrates, etc.), and (c) during development of individuals that start from a single cell and develop mechanisms that support different kinds of growth, development and learning, providing new mechanisms for processing information, at different stages of development, discussed briefly in the section on epigenesis, above. There can also be (d) new construction kits produced by a culture or ecosystem.

All such changes build on what was previously available, including bringing together separately evolved construction kits, for instance in symbiosis, sexual reproduction, or individual creative learning. One result has been evolution of toolkits that allow individuals to learn or construct any one of thousands of different human languages (whether spoken or signed) in the first few years of life. The evolutionary trajectory leading to human language capabilities may have gone from internal languages through collaborative actions then signed communication, then spoken communication, as argued in this presentation, where evidence is presented that constructing something new is more important than learning something old, in human language acquisition.

The history of computing since the first electronic computers in the 20th Century demonstrates some of the kinds of change that can arise when new construction kits are repeatedly based on previous kits. The technological changes are not merely changes of size, speed and memory capacity: there have been profound qualitative changes, in part because of the development of new layers of virtual machinery (Sloman (2013)) producing entirely new mechanisms based on old ones.

We now understand some of the key components and modes of composition (i.e. the basic components of a Computational Construction-Kit (CCK)) that constitute a platform on which all the human-designed layers of computation can be constructed. There are several different sorts of "basic" CCK that suffice to generate the forms of (discrete) computation so far studied. Those basic types include Turing machines, Post's production systems, Church's Lambda Calculus, and several more, each capable of generating the others. The Church-Turing thesis states that these construction kits generate all possible forms of computation, and there has been an enormous amount of research in computer science, and computer systems engineering, on forms of computation that can be built from such components. For more on this see: http://en.wikipedia.org/wiki/Church-Turing_thesis

However, it is not obvious that those discrete mechanisms suffice for all biological forms of information processing. The use of a wholly or partly chemical basis allows forms of computation that include discrete and continuous mechanisms that were essential for some forms of biological information processing, a topic that will be discussed in more detail elsewhere. Ganti (1971) shows how a chemical construction-kit supports forms of biological information processing that don't depend on external energy sources (a fact that's also true of battery-powered computers), and also supports growth and reproduction using internal mechanisms, which human-made computers cannot do.

There may be very different sorts of construction-kit that allow different sorts of information-processing (computation) to be supported, including some that we don't yet understand. In particular, the physical/chemical mechanisms that support the construction of both physical structures and information processing mechanisms in living organisms may have abilities not available in digital computers. Examples of human mathematical reasoning in geometry and topology that have so far resisted replication on computers are presented in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html

-- A new type of research project

Only a small subset of these biological processes and associated materials and mechanisms are well understood, though knowledge is increasing rapidly. As far as I know, very few of the derived construction kits have been identified and studied, and I am not aware of any systematic attempt to identify features of the FCK that explain the possibility of evolved biological DCKs. Researchers in fundamental physics or cosmology do not normally attempt to ensure that their theories explain not only known facts about physics and chemistry but also the many materials and process types that have been explored by natural selection and its products.

Schroedinger (1944) mentioned the need for a deep theory of the physical basis of life to explain such phenomena, though he could not have appreciated some of the requirements for increasingly sophisticated forms of information-processing, because, at the time he wrote, scientists and engineers had not learnt what we now know, after decades of computer systems engineering.

Curiously, although he mentioned the need to explain the occurrence of metamorphosis in organisms the example he mentioned was the transformation from a tadpole to a frog. He could have mentioned a far more spectacular type of example: the transformation from a caterpillar to a butterfly via an intermediate stage in which most of the caterpillar is transformed into a chemical soup within an outer case, from which the butterfly later emerges. See:
http://en.wikipedia.org/wiki/Pupa
http://en.wikipedia.org/wiki/Holometabolism

20th century biologists made some progress in understanding some of the achievements of the FCK in meeting physical and chemical requirements of various forms of life, though they used different terminology from mine, e.g. Haldane:
http://en.wikipedia.org/wiki/J._B._S._Haldane

However, the task can never be finished, since the process of construction of new derived construction kits may continue indefinitely, including new kits with components and modes of composition that allow production of more complex types of structure and more complex forms of behaviour in organisms.

That idea is familiar to computer scientists and computer systems engineers since many examples of new sorts of computational construction kit (new programming languages, new operating systems, new virtual machines) have been developed from old ones in the last half century, making possible new kinds of computing systems that could not previously be built from the original computing machinery, without introducing the new derived layers. This includes new virtual machines that are able to detect and record their own operations, a capability that is often essential for debugging and extending computing systems. For more on the importance of virtual machinery in extending what information-processing systems can do, and the properties they can have, see Sloman (2013).

In biological evolution over several billion years there have been vast numbers of new derived construction kits, making possible increasingly complex types of product, with increasingly complex behaviours. In some cases the components that are assembled are new complete organisms living together in cooperative or symbiotic forms, for example in termite cathedrals, beehives, ecosystems of various scales, and most recently human cultures and socio-economic systems. All such cases must make use of many DCKs of different kinds.

-- Construction-kits for biological information-processing

Layers of DCKs supported evolution not only of physical/chemical mechanisms, but later on a wide variety of information-processing capabilities, including perception, learning, motive formation, planning, decision making, and at a later stage mathematical discovery processes, by which individual life forms acquired mathematical competences that were useful in understanding sensory information, motive generation, plan construction, control of behaviour, and prediction. In some forms the mathematical discoveries were "compiled" into useful behaviours, e.g. use of negative feedback loops in control of temperature, osmotic pressure and other states, and use of geometric constraints by bees whose cooperative behaviours produce hexagonal cells in honeycombs.

Later still, construction kits used by evolution produced meta-cognitive mechanisms enabling individuals to notice and reflect on their own mathematical discoveries (enabling some of them to notice and remove flaws in their reasoning). In some cases those meta-cognitive capabilities allowed individuals to communicate their discoveries to others, discuss them, and organise them into complex highly structured bodies of shared knowledge, such as Euclid's *Elements*. Explaining how that could happen, and what it tells us about the nature of mathematics and biological/evolutionary foundations for mathematical knowledge is a long term goal of the Meta-Morphogenesis project, introduced above. A draft discussion is here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/befm-sloman.pdf

Many of these naturally occurring mathematical abilities have not yet been replicated in Artificial Intelligence systems or robots, unlike logical, arithmetical, and algebraic competences, which are well suited to computer modelling. Examples of topological reasoning not yet modelled in computers are presented here (and in other discussions linked from here):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html

It is not clear whether the difficulties in replicating such mathematical reasoning processes using AI techniques are due to the need for a kind of construction-kit that digital computers (e.g. Turing machines) cannot support, or due to our lack of imagination in using the capabilities of computers to replicate some biological forms of reasoning produced by evolution. Perhaps there are important forms of representation or types of information-processing architecture still to be discovered. There may also be properties of chemistry-based information-processing mechanisms that combine discrete and continuous interactions in ways that cannot be replicated exactly in discrete forms of computation. (This topic requires more detailed discussion.)

-- Representational blind spots of many scientists

Although I am not a physicist or mathematician and cannot follow all the details of writings of physicists, I think it is clear that most of the debates regarding what should go into a fundamental theory of matter ignore most of the biological demands on such a theory.

For example, presentations on physics make deep use of branches of mathematics concerned with numerical values, and the ways in which different measurable physical values do or do not co-vary, as expressed in (probabilistic or non-probabilistic) differential equations of various sorts. But the biological functions of complex physiological structures, especially structures that change in complexity, don't necessarily have those forms.

Biological mechanisms include: digestive mechanisms, mechanisms for transporting chemicals, mechanisms for detecting and repairing damage or infection, mechanisms for storing re-usable information about an extended structured environment, mechanisms for creating, storing and using complex percepts, thoughts, questions, values, preferences, desires, intentions and plans, including plans for cooperative behaviours, and mechanisms that transform themselves into new mechanisms with new structures and functions.

The forms of mathematics normally used by physicists are not necessarily useful for studying biological mechanisms. Logic, grammars and map-like representations are sometimes more appropriate, though I think little is actually known about the variety of forms of representation (i.e. encodings of information) used in human and animal minds and brains. We may need entirely new forms of mathematics for biology, and therefore for specifying what physicists need to explain.
Example: Many physicists, engineers and mathematicians who move into neuroscience assume that

states and processes in brains need to be expressed as collections of numerical measures and their derivatives plus equations linking them, a form of representation that is not necessarily best suited for the majority of mental contents, and probably not even well suited for chemical processes where structures form and interact with multiple changing geometrical and topological relationships -- one of the reasons for the invention of symbolic chemical notations (now being extended in computer models of changing interacting molecular structures). Information-processing mechanisms also often need to manipulate non-numerical structures.

-- Representing rewards, preferences, values
(Added 16 Feb 2015)

As an illustration of the previous point, many researchers assume that all decision making makes use of positive or negative scalar reward or utility values that are comparable across options between which decisions have to be made. But careful attention to consumer magazines, political debates, and the varieties of indecision that face humans in real life shows that reality is far more complex. For example, many preferences are expressed in rules about how to choose between certain options. Furthermore preferences can be highly sensitive to changes in context. A crude example is the change in preference for type of car after having children. An analysis of such examples led to the conclusion that "better" is a complex, polymorphic, logical concept with a rich structure that cannot be reduced to use of comparisons of numerical values (Sloman (1969),Sloman (1970).) Instead of a linear reward or utility metric real choices, for intelligent individuals, or for natural selection, will involve a complex partial ordering network, with "annotated" links between nodes. That philosophical conceptual analysis was later followed up by an attempt to design working models of intelligent agents with complex choices to be made under varying conditions, but the project merely scratched the surface, as reported in Beaudoin (1993), Beaudoin (1994), Wright et al. (1996), and Wright (1994) Despite all the sophistication of modern psychology and neuroscience, I don’t believe either has the resources to reveal the workings of brains in dealing with these matters. That includes being unable to identify the mechanisms leading to mathematical discoveries (new conjectures, new proofs, new counter-examples to conjectures), scientific theories, or new works of art and being unable to explain how they work. In part that’s because they lack sufficiently rich models of information processing (computation), and sufficiently deep methodologies for discovering what needs to be explained. Sloman (1969)

JUMP TO TABLE OF CONTENTS

Computational/Information-processing construction-kits

Since the mid 20th century we have been learning more and more about abstract construction-kits whose products are machines that can be used for increasingly complex tasks. Such construction kits include programming languages, operating systems, software development tools and environments, and network-technology that allows ever more complex information-processing machines to be constructed by combining simpler ones. A specially important, but poorly understood, feature of that history is growing use of construction-kits based on virtual machinery, discussed in some detail in Sloman (2013).

A complete account of the role of construction kits in biological evolution would need to include an explanation of how the fundamental construction kit (FCK) provided by the physical universe, as explained below, could be used later by evolution to produce a variety of types of virtual machinery as well as physical structures and mechanisms.

-- Infinite, or potentially infinite, generative power

A construction kit implicitly specifies a large, in some cases infinite, set of possibilities, though as an instance of the kit is constructed each addition of a new component or feature changes the set of possibilities accessible in later steps of that construction process.

For example, as you construct a sentence or phrase in a language, at each state in the construction there are alternative possible additions (not necessarily at the end) and each of those additions will alter the set of possible further additions consistent with the vocabulary and grammar of the language. When use of language is embedded in a larger activity, such as composing a poem, that context can modify the constraints that are relevant.

Chemistry does something like that for types of molecule, types of process involving molecular changes, and types of structure made of multiple molecules.

Quantum mechanics added important constraints to 19th century chemistry, including both the possibility of highly stable structures (resistant to thermal buffeting) and also locks and keys as in catalysis. All of that seems to be essential for life as we know it, and also for forms of information processing produced by evolution (mostly not yet charted).

One way of looking at research in fundamental physics is as a search for the construction kit that has the generative power to accommodate all the possible forms of matter, structure, process, causation, that exist in our universe.

However, physicists generally seek only to ensure that their construction kits are capable of accounting for phenomena observed in the physical sciences. Normally they do not assemble features of living matter, or processes of evolution, development, or learning, found in living organisms and try to ensure that their fundamental theories can account for those features also. There are notable exceptions such as Schroedinger (1944), but he, like most physicists who attempt to say something about physics and life (in my experience) ignored most of the details of life, including the variety of forms it can take, the variety of environments coped with, the different ways in which individual organisms cope, the ways in which products of evolution became more complex and more diverse over time, and the many kinds of information-processing and control in individuals, in colonies (e.g. ant colonies) and societies, and in ecosystems.

One of the issues some physicists have discussed is whether the formation of life from non-living matter requires violation of the second law of thermodynamics, because evolution increases the amount of order or structure in the physical matter on the planet. The standard answer is that the second law of thermodynamics is applicable only to closed systems, and the earth is not a closed system, since it is constantly affect by solar and other forms of radiation, asteroid impacts, and other external influences. I have tried to illustrate some of the ways in which pre-existing dispositions can harness external sources of energy to increase local structure in a short collection of thoughts on entropy, evolution, and construction-kits:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html

I suspect, but cannot (yet) demonstrate, that if cosmologists and other theoretical physicists attempted to take note of a wide range of biological phenomena (including the phenomena discussed in connection with the Meta-Morphogenesis project) they would find considerable explanatory gaps between current physical theories and the diversity of phenomena of life -- not because there is something about life that goes beyond what science can explain, but because we do not yet have a sufficiently rich theory of the constitution of the universe (or the Fundamental Construct Kit). I suspect that’s in part because of the forms of mathematics known to physicists. (The challenge of Anderson

(1972) is also relevant, as discussed below.)

Even if that is true it may take many years of research to find out what's missing from current physics. Collecting phenomena that need to be explained, and trying as hard as possible to construct detailed explanations of those phenomena is one way to make progress: it may help us to pin-point gaps in our theories and stimulate development of new more powerful theories, in something like the profound ways in which our understanding of possible forms of computation has been extended by attempts to put computation to new uses.

More importantly, collecting examples helps us assemble tests to be passed by future proposed theories: collections of possibilities that a deep physical theory needs to be able to explain.

Perhaps the most tendentious (controversial?) proposal here is that an expanded physical theory instead of being expressed mainly in terms of equations relating measures may need a formalism better suited to specification of a construction kit, perhaps sharing features of grammars, programming languages, partial orderings, topological relationships, architectural specifications, and the structural descriptions in chemistry -- all of which will need to make use of appropriate kinds of mathematics for drawing out implications of the theories, including explanations of possibilities, both observed and unobserved (e.g. future forms of intelligence?).

Types and levels of explanation of possibilities

Added 23 Nov 2014; Modified 22 Dec 2014; Major revisions 22 Jan 2015; 17 Feb 2015

[This section needs to be compressed.]

Suppose someone uses a meccano kit to construct a toy crane, with a jib that can be moved up and down by turning a handle, and a rotating platform on a fixed base, that allows the direction of the jib to be changed. What's the difference between explaining how that is possible and how it was done? First of all, if nobody actually builds such a crane then there is no actual crane-building to be explained: yet, insofar as the meccano kit makes cranes like that possible it makes sense to ask how it is possible. This has several types of answer, including answers at different levels of abstraction, with varying generality.

More generally, the question "How is it possible to create X using construction kit Y?" or, simply, "How is X possible?" has several types of answer, including answers at different levels of abstraction, with varying generality. I'll assume that a particular construction kit is referred to either explicitly or implicitly. The following is not intended to be an exhaustive survey of the possible types of answer.

1 Structural conformity: The first type of answer, structural conformity (grammaticality) merely identifies the parts and relationships between parts that are supported by the kit, showing that a crane of the sort in question could be composed of such parts arranged in such relationships. An architect's drawings for a building, specifying materials, components, and their spatial and functional relations would provide such an explanation of how a proposed building is possible, including, perhaps, answering questions about how the construction would make the building resistant to very high winds, or to earthquakes up to a specified strength. This can be compared with showing that a sentence is acceptable in a language with a well-defined grammar, by showing how the sentence would be parsed (analysed) in accordance with the grammar of that language. A parse tree (or graph) also shows how the sentence can be built up piecemeal from words and other grammatical units, by assembling various sub-structures and, using them to build larger structures. Compare using a chemical diagram to show how a collection of atoms can make up a particular molecule, e.g. the ring structure of C_6H_6 (Benzene).

Some structures are specified in terms of piece-wise relations, where the whole structure cannot possibly exist, because not all the relations can hold simultaneously, e.g. X is above Y, Y is above Z, Z is above X. It is even possible to depict such objects, e.g. pictures of impossible objects by Reutersvard, Escher, Penrose, and others. Some logicians and computer scientists have attempted to design languages in which specifications of impossible entities are necessarily syntactically ill-formed. This leads to impoverished languages that may have some practical uses (e.g. strongly typed programming languages) (Sloman (1971)).

2 Process possibility: The second type of answer demonstrates constructability by describing a sequence of spatial trajectories by which such a collection of parts could be assembled. This may include processes of assembly of temporary supports to hold parts in place before the connections have been made that make them self-supporting or before the final supporting structures have been built (as often happens in large engineering projects, such as bridge construction). Many different possible trajectories can lead to the same result. Describing (or demonstrating) any such trajectory explains both how that construction process is possible, and how the end result is possible. It may be possible to produce the end result in several different ways. In some cases a complex object has type 1 possibility although not type 2. For example, from a construction kit containing several rings it is possible to assemble a *pile* of three rings, but not possible to assemble a *chain* of three rings even though each of the parts of the chain is exactly like the parts of the pile.

3 Process Abstraction: Some possibilities described at a level of abstraction that ignores detailed routes through space, and covers **many** possible alternatives. For example, instead of specifying precise trajectories for parts as they are assembled, an explanation can specify the initial and final state of each trajectory, where those states may be shared by a vast, or even infinite collection, of different possible trajectories producing the same end state.

In some cases the possible trajectories for a moved component are all continuously deformable into one another (i.e. they are topologically equivalent): for example the many spatial routes by which a cup could be moved from a location where it rests on a table to a location where it rests on a saucer on the table, without leaving the volume of space above the table. Those trajectories form a continuum of possibilities that is too rich to be captured by a parametrised equation for a line, with a number of variables. If trajectories include leaving and entering the room via different doors or windows then the different possible trajectories will not all be continuously deformable into one another: there are different equivalence classes of trajectories sharing common start and end states.

The ability to abstract away from detailed differences between trajectories sharing start and end points, thereby implicitly recognizing invariant features of an infinite collection of possibilities, is an important aspect of animal intelligence that I don't think has been generally understood. Many researchers assume that intelligence involves finding *optimal* solutions. So they design mechanisms that search using an optimisation process, ignoring the possibility of mechanisms that can find sets of possible solutions (e.g. routes) initially considered as a class of *equivalent* options, leaving questions about optimal assembly to be settled later, if needed. These remarks are closely related to the origins of abilities to reason about geometry and topology, illustrated in these discussions:

These remarks are closely related to the origins of geometry and topology, illustrated in these discussions:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html

4 Grouping: Another form of abstraction is related to the difference between **1** and **2**. If there is a sub-sequence of assembly processes, whose order makes no difference to the end result, they can be grouped to form an unordered "composite" move, containing an unordered set of moves. If N components are moved from initial to final states in a sequence of N moves, and it makes no difference in what order they are moved, merely specifying the set of N possibilities without regard for order collapses N factorial sets of possible sequences into one composite move. If N is 10, that will collapse 3628800 different sequences into one. If each move can be represented only by start and end states, as in **3**, that will further reduce the space of alternatives.

Sometimes a subset of moves can be made in parallel. E.g. someone with two hands can move two or more objects at a time, in transferring a collection of items from one place to another. Parallelism is particularly important in many biological processes where different processes occurring in parallel constrain one another so as to ensure that instead of all the possible states that could occur by moving or assembling components separately, only those end states occur that are consistent with parallel constructions. Turing's discussion of parallel diffusion of two interacting chemicals in Turing (1952) provides examples.

This is important in epigenesis, since all forms of development from a single cell to a multi-celled structure depend on many mutually constraining processes occurring in parallel.

5 Iterative or recursive abstraction: Some process types involve unspecified numbers of steps, although each instance of the type has a definite number, for example a process of moving chairs by repeatedly carrying a chair to the next room until there are no chairs left to be carried, or building a tower from a collection of bricks, where the number of bricks can be varied. A specification that abstracts from the number can use a notion like "repeat until", or a recursive specification: a very old idea in mathematics, such as Euclid's algorithm for finding the highest common factor of two numbers. Production of such a generic specification can demonstrate possibilities inherent in a construction-kit in an extremely powerful and economical way. Many new forms of abstraction have been discovered by computer scientists developing programming languages, for operating not only on numbers but many other structures.

Evolution may also have discovered many, long before humans existed, by taking advantage of mathematical structures inherent in the construction-kits available and the trajectories by which parts can be assembled into larger wholes. This may be one of the ways in which evolution produced powerful new genomes, and re-usable genome components that allowed many different biological assembly processes to result from a single discovery at a high enough level of abstraction.

Some related abstractions may have resulted from processes by which details were removed from specifications in genomes and left to be provided by the context of development of individual organisms, including the physical or social environment. See the section on epigenesis below.

6 Self-assembly: If, unlike construction of a toy meccano crane or a sentence or a sorting process, the process to be explained is a self-assembly process, like many biological processes, then the explanation of how the assembly is possible will not merely have to specify trajectories through space by which the parts become assembled, but also
- What causes each of the movements (e.g. what manipulators are required)
- Where the energy required comes from (an internal store, or external supply?)
- Whether the process involves pre-specified information about the required end state or required steps, and if so what mechanisms can use that information to control the assembly process.
- How that prior information structure (e.g.specification of a goal state to be achieved, or plan specifying actions to be taken) came to exist, e.g. whether it was in the genome as a result of previous

evolutionary transitions, or whether it was constructed by some planning or problem-solving mechanism in an individual, or whether it was provided by a communication from an external source.
- How these abilities can be acquired or improved by learning or reasoning processes, or random variation (if they can).

Those are all examples of components of explanations of self-assembly processes. In biological reproduction, growth, repair, development, and learning there are far more subdivisions to be considered, some of them already studied piecemeal in a variety of disciplines. In the case of humans, and to a lesser extent other species there are many additional sub-cases involving construction kits for creating information structures and construction kits for creating information processing mechanisms of many kinds, including perception, learning, motive formation, motive comparison, intention formation, plan construction, plan execution, and many more. A subset of cases, with further references can be found in Sloman (2006).

The different answers to "How is it possible to construct this type of object" may be correct as far as they go, though some provide more detail than others. More subtle cases of explanations of possibility include differences between reproduction via egg-laying and reproduction via parturition. The latter allows a parent's influence to continue during development, as does teaching after birth/hatching.

-- Intentions and plans

None of the explanation-types above presupposes that the possibility being explained has ever been represented explicitly by the machines or organisms involved. Explaining the possibility of some structure or process that results from intentions or plans would require specifying pre-existing information about the end state and in some cases also intermediate states, namely information that existed before the process began -- information that can be used to control the process (e.g. intentions, instructions, or sub-goals, and preferences that help with selections between options).

Sometimes an explanation of possibility prior to construction is important for engineering projects where something new is proposed and critics believe that the object in question could not exist, or could not be brought into existence using available known materials and techniques. The designer might answer sceptical critics by combining answers of any of the above types, depending on the reasons for the scepticism.

JUMP TO TABLE OF CONTENTS

Alan Turing's Construction kits

Alan Turing showed in 1936 that a rather simple sort of machine, now known as a Turing machine, could be used to specify an infinite set of constructions with surprisingly rich mathematical features. The set of possibilities was infinite, because a Turing machine is defined to have an infinite (or indefinitely extendable) linear "tape" divided into discrete locations in which symbols can be inserted.

A feature of a Turing machine that is not in most other construction kits is that it can be set up and then started after which it will modify initial structures and build new ones, possibly indefinitely, though in some cases the machine will eventually halt.

Another type of construction kit with related properties is Conway's Game of Life, a self-propelling construction kit that creates changing patterns in 2D regular arrays. Stephen Wolfram has written a great deal about the diversity of constructions that can be explored using such cellular automata.

Here’s a video demo (one of many) on Youtube.

Neither a Turing machine nor a Conway game has any external sensors: once started they run according to their stored rules and the current (changing) state of the tape or grid-cells. In principle either of them could be attached to external sensors that could produce changes to the tape of a turing machine or the states of some of the cells in the Life array. However any such extension would significantly alter the powers of the machine, and theorems about what such a machine could or could not do might be invalidated.

Modern computers use a variant of the Turing machine idea where each computer has a finite memory but with the advantage of much more direct access between the central computer mechanism and the locations in the memory. (A von Neumann architecture.) Increasingly, computers have also been provided with a variety of external interfaces connected to sensors or motors so that while running they can acquire information (from keyboards, buttons, joy-sticks, mice, electronic piano keyboards, or network connections) and can also send signals to external devices. Theorems about disconnected Turing machines may not apply to machines with rich interfaces to an external environment. Many construction kits for learners are now available, enabling them to explore a wide variety of designs.

Not really self-propelling!

I described Turing machines and Game of Life machines as "self-propelling" because once you have set them up they run according to the general instructions they have and the initial configuration on the tape or in the array. But they are not really self-propelling: they have to be implemented in some sort of machine with an external power supply. In contrast, Ganti shows how the use of chemistry as a construction kit provides "self-propulsion" for living things, though every now and again the chemicals need to be replenished. A battery driven computer is a bit like that, but someone else has to make the battery.

Living things make and maintain themselves, at least after being given a kick-start by their parent or parents! They do need constant, or at least frequent, external inputs, but, for the simplest organisms, those are only chemicals in the environment, and energy either from chemicals or heat-energy via radiation, conduction or convection. John McCarthy pointed out in a conversation that some animals also use externally supplied mechanical energy, e.g. rising air currents used by birds. Unlike pollen-grains, spores, etc. propagated by wind or water, the birds use internal information-processing mechanisms to control how the wind energy is used, as does a human piloting a glider.

-- Beyond Turing machines: chemistry

Turing also explored other sorts of construction kits, including types of neural nets and extended versions of Turing machines with ’oracles’ added. In 1952, shortly before his death (in 1954), he published a paper in which he explored a type of pattern-forming construction kit in which two chemical substances can diffuse through the body of an expanding organism and interact strongly wherever they meet. He showed that that sort of construction kit could generate many of the types of surface physical structure observed on plants and animals.

One of the important differences between types of construction kit mentioned above is the difference between kits supporting only discrete changes (e.g. to a first approximation lego and meccano (ignoring variable length strings and variable angle joints) and kits supporting continuous variation, e.g. plasticine and mud (ignoring, for now, the discreteness at the molecular level).

One of the implications of such differences is how they affect abilities to search for solutions to problems. If only big changes in design are possible the precise change needed to solve a problem may be inaccessible (as I am sure many who have played with construction kits will have noticed). On the

other hand if the kit allows arbitrarily small changes it will, in principle, permit exhaustive searches in some sub-spaces. The exhaustiveness comes at the cost of a very much larger (infinite, or potentially infinite!) search-space. That feature could be useless, unless the space of requirements has a structure that allows approximate solutions to be useful. In that case a mixture of big jumps to get close to a good solution, followed by small jumps to home in on a (locally) optimal solution can be very fruitful: a technique that has been used by Artificial Intelligence researchers, called "simulated annealing". One of many online explanations is
http://www.theprojectspot.com/tutorial-post/simulated-annealing-algorithm-for-beginners/6

After I had written a first draft of this paper, Chris Scambler kindly drew my attention to a book Wagner (2014) claiming that the structure of the search space generated by the molecules making up the genome increases the chance of useful, approximate, solutions to important problems to be found with *relatively* little searching (compared with other search spaces), after which small random changes allow improvements to be found. I have not yet read the book but it seems to illustrate the importance for evolution of the types of construction-kit available. An interview with the author is online at https://www.youtube.com/watch?v=wyQgCMZdv6E.

As far as I know the book does not go into the uses of abstraction and the evolution of mathematical and meta-mathematical competences discussed here. Nevertheless, it seems to be an important (unwitting) contribution to the Meta-Morphogenesis project Sloman (2012+).

-- Using properties of a construction-kit to explain possibilities

A formal axiomatic system can be seen as an abstract construction kit with axioms and rules that support construction of proofs, ending in theorems. The theorems are formulae that can occur at the end of a proof using only axioms and inference rules in the system. The kit explains the possibility of some theorems based on the axioms and rules. The non-theorems of an axiomatic system are formulae for which no such proof exists. Proving that something is a non-theorem can be difficult, and requires a proof in a meta-system.

Likewise, a physical construction kit can be used to demonstrate that some complex physical objects can occur at the end of a construction process. In some cases there are objects that are describable but cannot occur in a construction using that kit: e.g. an object whose outer boundary is a surface that is everywhere curved, cannot be produced in a construction based on Lego bricks or a Meccano set, though one could occur in a construction based on plasticene, or soap-film.

-- Bounded and unbounded construction kits

A rectangular grid of squares combined with the single digit numbers, 0,1,..,9 (strictly numerals representing numbers) allows construction of a set of configurations in which numbers are inserted into the squares subject to various constraints, e.g. whether some squares can be left blank, or whether certain pairs of numbers can be adjacent, whether the same number can occur in more than one square. For a given grid and a given set of constraints here will be a finite set of possible configurations (although it may be a very large set).

If, in addition to insertion of a number, the "construction kit" allows extra empty rows or columns to be added to the grid, no matter how large it is, then the set of possible configurations becomes infinite. Many types of infinite construction kits have been investigated by mathematicians, logicians, linguists, computer scientists, musicians and other artists.

A more general investigation including chemistry-based construction kits for information processing systems is likely to range range over a far larger class of possible information processing systems than Turing machines (or digital computers) can, because of the mixture of discrete and continuous changes possible when molecules interact, e.g. moving together, moving apart, folding, twisting, but also locking and unlocking - using catalysts (Kauffman (1995)). But I don’t know whether anyone currently has a deep theory of the scope of chemistry-based information processing.

A key feature of the Meta-Morphogenesis project is investigation of construction kits that can produce new construction kits, including new construction kits able to produce construction kits that can produce new construction kits .... That’s why the process is Meta-Morphogenesis, not just Morphogenesis. (Connections between Meta-Morphogenesis and entropy are discussed in a related work-in-progress document on issues involving entropy changes in evolutionary trajectories http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html)

Conclusion

What sort of construction kit can support Meta-Morphogenesis?
(15 Feb 2015: Re-written and extended)

As I was finishing a first draft of this paper I found a useful survey by Evelyn Fox Keller of previous attempts to show how life and its products relate to the physical world, summarised (very briefly) here, which concluded that attempts so far have not been successful. However Keller ends with the suggestion that the traditional theory of dynamical systems is inadequate for dealing with constructive processes and needs to be expanded to include "objects, their internal properties, their construction, and their dynamics" i.e. a theory of *"Constructive dynamical systems"*. This paper outlines a project to do that and more: including branching layers of new *derived* construction kits produced by evolution, development and other processes. My suggestion is that the physical world provides a very powerful (mostly chemical) fundamental construction kit that, together with natural selection processes and processes within individuals as they develop, produced an enormous variety of organisms on this planet, based on additional derived construction kits (DCKs), including concrete, abstract and hybrid construction kits, and most, recently, new sorts of construction kit used as toys or engineering resources.

This use of the idea of construction kits is, as far as I know, new to philosophy of mathematics, philosophy of science, philosophy of biology, philosophy of mind and metaphysics, although special cases are familiar. But the idea is still at an early stage of development and there are probably many more distinctions to be made, and a need for a more formal, mathematical presentation of properties of and relationships between construction kits, including the ways in which new derived construction kits can be related to their predecessors and their successors. The many new types of computer-based *virtual* machinery produced (intentionally and unintentionally) by human engineers since around 1950 providing examples of non-reductive supervenience (as explained in Sloman (2013)) are useful as relatively simple examples to be compared with far more complex products of evolution.

The general idea of a construction kit used in this paper is, as far as I know, new to philosophy of mathematics, philosophy of science, philosophy of mind and metaphysics, although special cases are familiar. But the idea is still at an early stage of development and there are probably many more distinctions to be made, and a need for a more formal, mathematical presentation of properties of and relationships between construction kits/, including the ways in which new derived construction kits can be related to their predecessors and their successors. The many new types of computer-based virtual machinery produced (intentionally and unintentionally) by human engineers since around 1950 providing examples of non-reductive supervenience (as explained in Sloman (2013)) are useful as relatively simple examples to be compared with far more complex products of evolution.

In Esfeld (in press) a distinction is made between two "principled" options for the relationship between the basic constituents of the world and their consequences. In the "Humean" option there is nothing but the distribution of structures and processes over space and time, though there may be some empirically discernible patterns in that distribution. The second option is "modal realism", or "dispositionalism", according to which there is something about the primitive stuff and its role in space-time that constrains what can and cannot exist, and what types of process can or cannot occur. This paper supports a "multi-layer" version of the modal realist option (developing ideas in (Sloman 1962), (Sloman 1996) and (Sloman 2013)).

I conjecture that a more complete development of this form of modal realism can contribute to answering the problem posed in Anderson’s famous paper Anderson (1972), namely how we should understand the relationships between different levels of complexity in the universe (or in scientific theories). The reductionist alternative claims that when the physics of elementary particles (or some other fundamental physical level) has been fully understood, everything else in the universe can be explained in terms of mathematically derivable consequences of the basic physics. Anderson contrasts this with the anti-reductionist view that different levels of complexity in the universe require "entirely new laws, concepts and generalisations" so that, for example, biology is not applied chemistry and psychology is not applied biology. He writes: "Surely there are more levels of organization between human ethology and DNA than there are between DNA and quantum electrodynamics, and each level can require a whole new conceptual structure". I would only add that the structural levels are not merely in the concepts used by scientists, but in the world.

We still have much to learn about the powers of the fundamental construction kit (FCK), including: (i) the details of how those powers came to be used for life on earth, (ii) which sorts of derived construction kit (DCK) were required in order to make more complex life forms possible, (iii) how those construction kits support "blind" mathematical discovery by evolution, mathematical competences in humans and other animals and eventually meta-mathematical competences, then meta-meta-mathematical competences, at least in humans, (iv) what sorts of potential the FCK has that have not yet been realised, (v) whether and how some version of the FCK could be used to extend the intelligence of current robots, and (vi) whether currently used Turing-equivalent forms of computation have at least the same information-processing potentialities (e.g. abilities to support all the biological information-processing mechanisms and architectures), and (vii) if those forms of computation lack the potential, then how are biological forms of information processing different? Don’t expect complete answers soon.

In future, physicists wishing to show the superiority of their theories, should attempt to demonstrate mathematically and experimentally that they can explain more of the potential of the FCK to support varieties of construction kit required for, and produced by, biological evolution than rival theories can. Will that be cheaper than building bigger better colliders? Will it be harder?
Here’s a cartoon teasing particle physicists:
http://www.smbc-comics.com/?id=3554
Thanks to Tanya Goldhaber for the link.

JUMP TO TABLE OF CONTENTS

REFERENCES

- http://robotics.cs.tamu.edu/dshell/cs689/papers/anderson72more_is_different.pdf Philip W. Anderson, (1972), More is different, *Science, New Series*, 177, 4047, pp. 393--396,
- http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#16
  L. Beaudoin A. Sloman A (1993) A study of motive processing and attention. In:
  Sloman A, Hogg D, Humphreys G, Partridge D, Ramsay A (eds) *Prospects for Artificial Intelligence,*
  IOS Press, Amsterdam, pp 229-238.
- http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#38
  L. Beaudoin (1994) *Goal processing in autonomous agents.*
  PhD thesis, School of Computer Science, The University of Birmingham,
- Graham Bell, 2008, *Selection The Mechanism of Evolution,* 2nd Edn, OUP.
- V. Braitenberg, 1984, *Vehicles: Experiments in Synthetic Psychology,* The MIT Press,
- Natural and artificial meta-configured altricial information-processing systems
  Jackie Chappell and Aaron Sloman
  International Journal of Unconventional Computing, 3, 3, 2007, pp. 211--239,
  http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717
- Juliet Coates, Umm-E-Aiman and Benedicte Charrier, Dec, 2014,
  Understanding green multicellularity: do seaweeds hold the key? in *Frontiers in Plant Science,*
  doi: 10.3389/fpls.2014.00737,
  http://journal.frontiersin.org/Journal/10.3389/fpls.2014.00737/pdf
- http://arxiv.org/abs/1411.7545
  Michael Esfeld, Dustin Lazarovici, Vincent Lam and Mario Hubert,
  The Physics and Metaphysics of Primitive Stuff, *British Journal for the Philosophy of Science,*
  (Forthcoming)
- Gibson J.J (1966) The Senses Considered as Perceptual Systems. Houghton Mifflin, Boston
- Gibson J.J. (1979)
  The Ecological Approach to Visual Perception. Houghton Mifflin, Boston, MA
- Eva Jablonka and Marion J. Lamb, 2005, *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life,* MIT Press
- Hameroff S, Penrose R (2014) Consciousness in the universe: A review of the orch or theory. Physics of Life Reviews 11(1):39 - 78, DOI http://dx.doi.org/10.1016/j.plrev.2013.08.002
- Immanuel Kant, *Critique of Pure Reason,*
  1781, Translated (1929) by Norman Kemp Smith, London, Macmillan,
- Annette Karmiloff-Smith
  Beyond Modularity,
  A Developmental Perspective on Cognitive Science,
  MIT Press (1992) --Informally reviewed in
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html

- Stuart Kauffman (1995) *At home in the universe: The search for laws of complexity*. Penguin Books, London

- Minsky ML (1987) *The Society of Mind.* William Heinemann Ltd., London

- Minsky ML (2006) *The Emotion Machine.* Pantheon, New York

- Roger Penrose (1994) *Shadows of the mind: A Search for the Missing Science of Consciousness.* OUP, Oxford

- Ernst Schroedinger (1944) *What is life?* CUP, Cambridge

- http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#706
  Aaron Sloman (1962)
  DPhil Thesis: Knowing and Understanding:
  Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth. University of Oxford.

- http://www.cs.bham.ac.uk/research/cogaff/papers.html#1969-02
  Aaron Sloman, 1969, How to derive "Better" from "is", in
  American Phil. Quarterly, 6, pp. 43--52,

- http://www.cs.bham.ac.uk/research/projects/cogaff/05.html#200508
  Sloman A (1970) "Ought" and "better". *Mind* LXXIX(315):385-394,

- http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200304
  Sloman A (1971) Tarski, Frege and the Liar Paradox. *Philosophy* 46(176):133-147,

- http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2
  Aaron Sloman (1978)
  'What are the Aims of Science?' Chapter 2 of *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind*
  Whole book: http://www.cs.bham.ac.uk/research/projects/cogaff/crp/

- http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#18
  Aaron Sloman, 1993, The mind as a control system, in *Philosophy and the Cognitive Sciences*,
  Eds. C. Hookway and D. Peterson, CUP pp. 69--110,

- http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15
  Aaron Sloman, 1996 Actual Possibilities, in
  Principles of Knowledge Representation and Reasoning: Proc. 5th Int. Conf. (KR '96),
  Eds. L.C. Aiello and S.C. Shapiro, Morgan Kaufmann Publishers, Boston, MA, pp. 627--638,

- http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#62
  A. Sloman, (2000) Interacting trajectories in design space and niche space:
  A philosopher speculates about evolution, in
  Parallel Problem Solving from Nature -- PPSN VI,
  Eds. M.Schoenauer, et al., Springer-Verlag, LNCS No 1917, pp. 3--16.

- Aaron Sloman (2003), *The Cognition and Affect Project: Architectures, Architecture-Schemas, And The New Science of Mind.,*
  School of Computer Science, University of Birmingham, UK,
  http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200307

- http://www.cs.bham.ac.uk/research/projects/cogaff/misc/fully-deliberative.html
  Aaron Sloman, (2006),
  Requirements for a Fully Deliberative Architecture (Or component of an architecture),
  Research Note, COSY-DP-0604,
  School of Computer Science, University of Birmingham, UK,

- http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#sab
  Aaron Sloman, 2010,
  How Virtual Machinery Can Bridge the "Explanatory Gap", In Natural and Artificial Systems,
  Proceedings SAB 2010, LNAI 6226, Eds. S. Doncieux and et al., Springer, pp. 13--24.

- http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html
  Aaron Sloman, Current website for the Meta-Morphogenesis project, originally set up in 2012, based on this paper:
  Aaron Sloman, (2013), Virtual Machinery and Evolution of Mind (Part 3): Meta-Morphogenesis: Evolution of Information-Processing Machinery, in *Alan Turing - His Work and Impact*,
  Eds. S. B. Cooper and J. van Leeuwen, Elsevier, Amsterdam, pp. 849-856,
  http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d
  (Full contents of the book)

- http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html
  Aaron Sloman (2013, Work in Progress)
  Virtual Machine Functionalism: (The only form of functionalism worth taking seriously in Philosophy of Mind)
  Discussion Note, University of Birmingham

- Stephen P. Stich (1981)
  Review: *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind* (1978) by Aaron Sloman, in *The Philosophical Review*, Vol. 90, No. 2 (Apr., 1981), pp. 300-307
  Available here with the author’s permission:
  http://www.cs.bham.ac.uk/research/projects/cogaff/crp/stich-review-crp.html

- A. Tarski, (1944), The semantic conception of truth, in *Philosophy and Phenomenological Research*, vol 4, no 3, pp. 341--376,

- A. M. Turing, (1952), The Chemical Basis Of Morphogenesis, *Phil. Trans. R. Soc. London*, B 237, 237, pp. 37--72,

- C. H. Waddington, *The Strategy of the Genes*, 1957, MacMillan.

- Andreas Wagner, 2014, *Arrival of the Fittest: Solving Evolution’s Greatest Puzzle*
  Published by: Oneworld Publications,

- http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#2
  Wright I, Sloman A, Beaudoin L (1996)
  Towards a design-based analysis of emotional episodes. In *Philosophy Psychiatry and Psychology* Vol3, No2, 101-126

- http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#22
  Wright I (1977) *Emotional agents.*
  PhD thesis, School of Computer Science, The University of Birmingham,