
Under Construction: Comments Welcome

New title: 8 Apr 2016

The previous title: Construction kits for biological evolution

Was used for my chapter in the Springer 2017 book

The Incomputable: Journeys Beyond the Turing Barrier (see below).

This document is work in progress, extending that chapter.

NOTE (13 Jun 2017) I have just realised that this work implies that evolution uses a form of **concrete abduction** loosely analogous to reasoning by abduction.

The document will be revised accordingly.

Construction kits for evolving life (Including evolving minds and mathematical abilities.)

The scientific/metaphysical explanatory role of construction kits:

**fundamental and derived kits,
concrete, abstract and hybrid kits,
meta-construction kits.**

(All with deep mathematical properties, creating products with mathematical properties.)

Aaron Sloman

School of Computer Science, University of Birmingham, UK

<http://www.cs.bham.ac.uk/~axs>

Still growing/changing: Saved copies will grow out of date.

This document is available in two formats: HTML and PDF

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

PDF: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.pdf>

(The PDF version may not always include the latest changes.)

This is part of: The Turing-inspired Meta-Morphogenesis (M-M) project,


(Or Self-Informing Universe (SIU) project)

which asks:

**How can a cloud of dust give birth to a planet
full of living things as diverse as life on Earth?**

(Including mathematicians like Archimedes, Euclid, Zeno, and many others.)

How can a protoplanetary dust cloud give birth to Earth life? What features of the physical universe make that possible?



[[NASA artist's impression of a protoplanetary disk, from Wikimedia](#)]

[Jump to CONTENTS](#)

BACKGROUND NOTES

This is part of the Meta-Morphogenesis (M-M) project.

Additional topics are included or linked at the main M-M web site:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

The Beginning: Two book chapters

Ideas about Meta-Morphogenesis in Biological evolution were first presented in a paper written in 2011, as my invited contribution [[Sloman 2013b](#)], to part 4 of [[Cooper and van Leeuwen 2013](#)], commenting on Turing's paper on The Chemical basis of Morphogenesis [[Turing 1952](#)]. That paper introduced the Meta-Morphogenesis project and the conjecture that Turing would have worked on such a project if he had not died two years after publishing the [morphogenesis paper](#).

Since then, several published and unpublished papers on [the Birmingham CogAff website](#) have extended the ideas in several directions, one of the most important being the theory of evolved construction kits, presented below, but still in its infancy. Another important theme is the creativity of evolution. Not only are the biological mechanisms of evolution enormously creative, they are also indirectly responsible for all the other types of creativity on this planet, e.g. human creativity.

Some of the ideas concerning creativity of evolution are in this incomplete paper on creativity:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/creativity.html>

(work in progress). See also [[Sloman 2012](#)].

The main idea behind the label "Meta-Morphogenesis" was that some of the changes produced by the mechanisms of evolution (an important type of morphogenesis) extend the mechanisms of evolution (hence the 'Meta-' in 'Meta-morphogenesis').

The most significant products of evolution able to feed back into mechanisms of evolution include both many physical/chemical products of evolution and also [evolved forms of information processing](#).

Most of the advances in studies of evolution have been concerned with the physical and chemical structures, along with the physical forms of organisms, the environments with which they have coped, and aspects of their behavioural competences inferred from observations of existing species, and speculative inferences from fossil records and archeological evidence.

But the ancient forms of information processing do not leave fossil records, although some of their products do. And I do not know of any systematic attempt to use varied forms of direct and indirect evidence to infer ancient forms and mechanisms of biological information processing.

In contrast, inspired by Turing's work, the main aim of the M-M project, from the beginning, was to identify changes in types of information processing produced by evolution, including new forms of information-processing that directly or indirectly altered the scope of biological evolution (hence the "Meta-").

Two important examples of such changes were the development of sexual reproduction, and later on the development of mate-selection mechanisms. However there must have been many more changes concerned with uses of information of many kinds for many purposes, including collaborative uses of information, e.g. in swarming and trail-following behaviours.

After I had contributed four papers to the [Turing Centenary Volume](#), including the paper in Part IV introducing the Meta-Morphogenesis project without mentioning construction kits, Barry Cooper and Mariya Soskova invited me to give a talk at a Workshop on "The Incomputable", held at the Kavli Centre in June 2012, at which I presented some of the ideas about the M-M project. They later invited me to contribute a chapter for a follow-up book: The focus on construction kits came out of that invitation:

"This new book will aim to address the theme of incomputability with a wide readership in mind, filling a real gap in the literature. Our choice of potential contributors prioritises leading researchers who can write, and are able to contribute something for both experts and non-experts. We see the book as uniquely focusing on this neglected aspect of the Turing legacy, and sharing more generally some appreciation of the beauty and fundamental importance of contemporary research at the computability/incomputable interface."

Because of my special interest (since my thesis defending Kantian philosophy of mathematics [1962](#)) in the nature of ancient mathematical forms of reasoning and discovery (e.g. the amazing deep ancient discoveries in geometry and topology reported by Archimedes, Euclid, Zeno and others) with features that had so far resisted replication or modelling in AI theorem provers, and which did not seem to be explicable in terms of current theories about brain mechanisms, I started writing a paper presenting examples of types of mathematical reasoning and discovery that were not yet replicated in AI, intending to propose a strategy for investigating whether the gaps were simply due to limitations in current ideas about automated reasoning, or an indication of some more fundamental, previously unnoticed, problem about what current AI models of computation can achieve.

My ideas about possible gaps were different from the ideas of AI critics such as Searle (1981) and Penrose ([1989](#)) and ([1994](#)), both of whom I had criticised in [Sloman\(1992\)](#) and other papers.

So, in the spirit of the M-M project, I proposed collecting examples of evolutionary transitions in biological information processing that might eventually have led to evolution of ancient mathematical minds. The idea that construction kits would have to play a role was triggered by a talk given by Birmingham Biologist [Juliet Coates](#) on evolution of biological toolkits that were necessary for the development of the earliest plants, e.g. enabling them to produce structures that permitted vertical growth upwards from a supporting medium. (See [[Coates, et al. 2014](#)].)

This made me realise that there was a very general notion missing from the Darwin/Wallace theory of evolution in all the versions that I had encountered, and the plant toolkit was a special case.

The missing notion was that evolution did not merely produce all the different types of organism on the planet, with all the different types of body-part with different functions, and different observed behaviours. Evolution must also have produced many different sorts of construction kit, including construction kits for producing and extending information processing mechanisms, some of which will have required new body parts, in the same way as human engineered information processing mechanisms use different sorts of physical mechanism, while other construction kits may have produced more abstract mechanisms, in something like the way human information engineering has produced both new programming languages, new programs, new operating systems, new software tools for use in building, testing and maintaining new packages, and so on.

NB. None of this implies that the biological information processing systems were Turing machines, or were capable of being implemented on Turing machines, or even that all the forms of information processing were [digital](#), using only discrete physical structures and processes. Human engineers have used both digital and analog (continuously varying) mechanisms for centuries, and I see no reason for ruling out a combination. Turing's [morphogenesis paper](#) indicates that he was interested in forms of continuous variation that could produce both continuous and discrete changes, suggesting that he was more open minded than some of his admirers.

These more abstract construction kits and their products are as important for the M-M project as the construction kits for building and assembling new physical/chemical components. This is something Schrödinger appreciated in [[1944](#)], before the development of hardware and software computing technologies had begun.

The whole process of creating designs for organisms, parts of organisms, and construction kits for instantiating those designs, must have started with a "Fundamental" construction kit (FCK) with the potential to generate all the other construction kits. That FCK must have been provided by fundamental physical features of the universe, long before there was any life on this planet (or anywhere else).

This implied that the M-M theory had to be extended to accommodate construction kits for generating construction kits -- an idea that should be familiar to anyone with experience of designing and building complex software systems. (My own experience included contributions to Poplog, and the SimAgent toolkit, built on Poplog [[Sloman 1996b](#)].)

In order to identify missing components in our current toolkits for building information processing systems I began collecting examples of types of mathematical discovery that did not fit current AI reasoning systems, and trying to categorise them in different types and subtypes.

That process converted my contribution to the *Incomputable* book into a first draft investigation of types of construction kit required by biological evolution. I hoped the project would shed new light on biological construction kits for information processing mechanisms that might have played a role in the evolution of mathematical discovery and reasoning mechanisms that were already producing rich results several thousand years ago, long before the invention/discovery of modern logic, algebra, set theory, formal systems, proof theory, etc.

An initial progress report, [[Sloman 2017a](#)], entitled "Construction Kits for Biological Evolution", presenting ideas developed between 2014 and 2016, was eventually published by Springer in 2017, in [The Incomputable, \(Cooper and Soskova 2017\)](#), although, sadly, Barry Cooper whose encouragement was crucial to this project, did not live to see it.

(The book chapter was initially frozen in Dec 2015, then modified in September and December 2016, especially undoing many erroneous changes by Springer copy editors. (See my rant against copy editors [here](#).)

This paper continues the work on construction-kits, perhaps the single most important part of the M-M project, for the time being. Several parts of the book chapter have been extended and partly rewritten. This version is still changing, and the externally visible online version will be changed from time to time. Please store links, not copies, as copies will become out of date.

Evolved Compositionality (Added 30 Nov 2018, after SYCO-1 Workshop. Modified 12 Feb 2019)

Most philosophies of learning and science assume there's a fixed type of world to be understood by a learner or a community of scientists. Biological evolution, however, in collaboration with its products, is metaphysically creative, and constantly extends the mathematical diversity and complexity of the world to be understood, and constantly extends the types of learners with new powers. Its more advanced (recently evolved) learners have multi-stage genomes that extend powers within a learner in different ways at different stages (during epigenesis), by extending the powers produced by the earlier stages [and their environments](#), which can differ across generations and across geographical locations. This process constantly adds [new compositional powers](#) (new types of compositionality) making new, more complex step-changes available to natural selection. Most will probably be fatal, but enough are viable for to have a major impact on the generative power of biological evolution. (This all needs to be documented in far more detail -- unless someone has already done that.)

So evolution extends genomes, which get changed by evolution, and new multi-layered genomes extend the ways in which individuals change themselves as they as they get changed by the environment.

Moreover individuals with new genomes can change the environment in new ways, producing new types of selection pressure.

Sexual reproduction can combine different evolved structures with different histories making step-changes to the structural variety and potential in the gene-pool.

The contrast between the powers of [genetic algorithms](#) (GAs) and [Genetic programming](#) (GP) illustrates some of the differences discussed here. [The GA/GP contrast is controversial, however, and there are likely to be differences between the impact in computer-based experiments on artificial evolution and the impact of the mechanisms described here on biological evolution. For more information and references, see https://en.wikipedia.org/wiki/Genetic_programming]

Together, all these processes combined with non biological external influences can use the compositionality of the genome [to produce step changes at different levels of abstraction into new regions of the space of possibilities](#). This continually produces not only new designs for individuals, but also extends the space of available reproductive and developmental trajectories, including continually adding new, more abstract and general levels of compositionality, allowing previous products of evolution to be combined in new ways. Hence the label: The Meta-Configured genome -- explained more fully [below](#).

New capabilities can combine with new forms of motivation to produce new behavioural capabilities able to contribute to achieve new types of goal. Some of these processes also requires new forms of motivation during development, for example new forms of play, and in some cases play fighting.

Eventually these processes produced minds able to make deep mathematical discoveries and to use the results in increasingly sophisticated scientific theories and engineering applications. But we still don't know how to replicate those powers in human-designed machines, despite all the advances in AI.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html>

A more detailed discussion of the role of compositionality in evolution is in:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.pdf>

[Contrast: "Ontogeny recapitulates phylogeny" (Haeckel), a much simpler idea.]

Updated: 29 Nov 2018 ([Compositionality](#)); 13 Jan 2019; 12 Feb 2019

14 Jan 2018 (Memo functions and genome size); 1 Jun 2018; 11 Nov 2018;

May 2017; 3 Oct 2017(Minor corrections);

10 Oct 2016 (including correcting "kit's" to "kits!"); 10 Feb 2017;

3 Feb 2016; 10 Feb 2016; 20 Feb 2016; 23 Feb 2016; 19 Mar 2016; 30 Jun 2016; 27 Dec 2016

Original version: December 2014

[Jump to CONTENTS](#)

Other versions and related papers.

This document is occasionally copied to slideshare.net, making it available in flash format. That version will not necessarily be updated whenever the html/pdf versions are. (Not all the links work in the slideshare version.) Slideshare version last updated 10 Oct 2016.

See <http://www.slideshare.net/asloman/construction-kits>

(Note: Slideshare now does not allow updating, a dreadful change of policy.)

Closely related online papers:
(Many also have PDF versions)

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

Overview of the Meta-Morphogenesis project, and some of its history.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html>

**Alan Turing's 1938 thoughts on intuition vs ingenuity in mathematical reasoning
Did he unwittingly re-discover key aspects of Kant's philosophy of mathematics, illustrated in**
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html>

**Why can't (current) machines reason like Euclid or even human toddlers?
(And many other intelligent animals)**

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>

**Virtual Machine Functionalism (VMF)
(The only form of functionalism worth taking seriously
in Philosophy of Mind and theories of Consciousness)**

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/maths-multiple-foundations.html>

**Multiple Foundations For Mathematics
Neo-Kantian (epistemic/cognitive) foundations,
Mathematical foundations, Biological/evolutionary foundations
Cosmological/physical/chemical foundations, Metaphysical/Ontological foundations
Multi-layered foundations, others ???**

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html>

**Using construction kits to explain possibilities
(Defending the scientific role for explanations of possibilities, not just laws.)**

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/creativity.html>

The Creative Universe
(Early draft, begun March 2016)

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

The Birmingham Cognition and Affect project
(begun 1991, extending ideas developed earlier while I was at Sussex University).

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html>

**Tentative non-mathematical thoughts on entropy, evolution, and construction-kits
(Entropy, Evolution and Lionel Penrose's Droguli)**

Ongoing work on fundamental and derived construction kits is summarised below.

A partial index of papers and discussion notes on this web site is in
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html>

A separate "background" document on Evelyn Fox Keller (1 Mar 2015)

A few notes on [Evelyn Fox Keller's](#) papers on Organisms, Machines, and Thunderstorms: A History of Self-Organization, in *Historical Studies in the Natural Sciences*, Vol. 38, No. 1 (Winter 2008), pp. 45-75 and Vol. 39, No. 1 (Winter 2009), pp. 1-31
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/keller-org.html>

NOTE 16 Jan 2018

Related work in Biosemiotics

Alexei Sharov drew my attention to deep, closely related work by the [Biosemiotics](#) research community, e.g.:

<http://www.biosemiotics.org/biosemiotics-introduction/>

<https://en.wikipedia.org/wiki/Biosemiotics> <https://en.wikipedia.org/wiki/Zoosemiotics>

I shall later try to write something about the connections, and will add links on this site.

Contents

[BACKGROUND NOTES \(above\)](#)

[Evolved Compositionality \(Above\)](#)

[A separate "background" document on Evelyn Fox Keller \(Above\)](#)

[ABSTRACT: The need for construction kits \(Below\)](#)

[Philosophical background: What is science? Beyond Popper and Lakatos](#)

[Note on "Making Possible":](#)

1 [A Brief History of Construction-kits](#)

[Beyond supervenience to richer forms of support](#)

[A corollary: a tree/network of evolved construction kits](#)

[The role\(s\) of information in life and its evolution](#)

2 [Fundamental and Derived Construction Kits \(FCK, DCKs\)](#)

[SMBC comic-strip comment on "fundamentality"](#)

2.1 [Combinatorics of construction processes](#)

[Abstraction via parametrization](#)

[New tools, including virtual machinery](#)

[NOTE on vision \(Added 3 Sep 2017\)](#)

[Combinatorics of biological constructions](#)

[Comparison with use of Memo-functions in computing](#)

[Storing solutions vs storing information about solutions](#)

[Negative memoization](#)

[NOTE: making faster or easier vs making possible:](#)

[Figure FCK: The Fundamental Construction Kit](#)

[Figure DCK: Derived Construction Kits](#)

2.2 [Construction Kit Ontologies](#)

2.3 [Construction kits built during development \(epigenesis\)](#)

[Figure EVO-DEVO](#)

2.4 [The variety of biological construction kits](#)

[NOTE: Refactoring of designs \(Added 9 Oct 2016\)](#)

- [2.5 Increasingly varied mathematical structures](#)
 - [IJCAI 2017 presentation \(Added 29 Aug 2017\)](#)
 - [2.6 Thermodynamic issues](#)
 - [2.7 Scaffolding in construction kits](#)
 - [2.8 Biological construction kits](#)
 - [2.9 Cognition echoes reality](#)
 - [\(INCOMPLETE DRAFT. Added 13 Apr 2016\)](#)
 - [2.10 Goodwin's "Laws of Form": Evolution as a form-maker and form-user](#)
 - [\(DRAFT Added: 18 Apr 2016\)](#)
 - [3 Concrete \(physical\), abstract and hybrid construction kits](#)
 - [3.1 Kits providing external sensors and motors](#)
 - [3.2 Mechanisms for storing, transforming and using information](#)
 - [3.3 Mechanisms for controlling position, motion and timing](#)
 - [3.4 Combining construction kits](#)
 - [3.5 Combining abstract construction kits](#)
 - [4 Construction kits generate possibilities and impossibilities](#)
 - [4.1 Construction kits for making information-users](#)
 - [4.2 Different roles for information](#)
 - [Figure Evol \(Evolutionary transitions\)](#)
 - [4.3 Motivational mechanisms](#)
 - [5 Mathematics: Some constructions exclude or necessitate others](#)
 - [5.1 Proof-like features of evolution](#)
 - [5.2 Euclid's construction kit](#)
 - [5.3 Mathematical discoveries based on exploring construction kits](#)
 - [5.4 Evolution's \(blind\) mathematical discoveries](#)
 - [Dana Scott's new route to Euclidean geometry](#)
 - [6 Varieties of Derived Construction Kit](#)
 - [6.1 A new type of research project](#)
 - [6.2 Construction-kits for biological information processing](#)
 - [6.3 Representational blind spots of many scientists](#)
 - [6.4 Representing rewards, preferences, values](#)
 - [7 Computational/Information-processing construction-kits](#)
 - [7.1 Infinite, or potentially infinite, generative power](#)
 - [8 Types and levels of explanation of possibilities](#)
 - [9 Alan Turing's Construction kits](#)
 - [9.1 Beyond Turing machines: chemistry](#)
 - [9.2 Using properties of a construction-kit to explain possibilities](#)
 - [9.3 Bounded and unbounded construction kits](#)
 - [9.4 Towers vs forests \(Draft: 2 Feb 2016\)](#)
 - [9 Alan Turing's Construction kits](#)
 - [10 Conclusion: Construction kits for Meta-Morphogenesis](#)
 - [11 End Note \(Turing's letter to Ashby\)](#)
 - [12 Note on Barry Cooper](#)
- [FOOTNOTES](#)
- [REFERENCES](#)
-

[Jump to CONTENTS](#)

ABSTRACT: The need for construction kits

This is part of the Turing-inspired Meta-Morphogenesis project, which aims to identify transitions in information processing since the earliest proto-organisms, partly in order to provide new understanding of varieties of biological intelligence, including the mathematical intelligence that produced Euclid's *Elements*, without which a great deal of human science and engineering would have been impossible. (Explaining evolution of mathematicians is much harder than explaining evolution of consciousness, since being able to make mathematical discoveries requires common forms of consciousness with some additional features!)

Evolutionary transitions depend on the availability of "construction kits", including the initial "Fundamental Construction Kit" (FCK) based on physics and chemistry, and "Derived Construction Kits" (DCKs) produced by combinations of physical processes (e.g. lava-flows, geothermal activity) and biological evolution, development, learning and culture. Some are meta-construction kits: construction kits for building new construction kits.

Some construction kits used and in many cases also designed by humans (e.g. Lego, Meccano, plasticine, sand, piles of rocks) 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. Concrete and abstract construction kits can be combined to form *hybrid* kits, e.g. games like tennis or croquet, which involve physical objects, players, and rules. There are also *meta-construction kits*: able to create, modify or combine construction kits. (This list of types of construction kit is provisional, and likely to be extended, including new sub-divisions of some of the types.)

Construction kits are *generative*: they explain sets of possible construction processes, and possible products, often with mathematical properties and limitations that are mathematical consequences of properties of the kit and its environment (including general properties of space-time).

Evolution and individual development processes can both make new construction kits possible. Study of the FCK and DCKs can lead us to new answers to old questions, e.g. about the nature of mathematics, language, mind, science, and life, exposing deep connections between science and metaphysics, and helping to explain the extraordinary creativity of biological evolution. (Discussed further in a separate document [on creativity](#).)

Products of the construction kits are initially increasingly complex **physical** structures/mechanisms. Later products include increasingly complex **virtual** machines.

Philosophers and scientists were mostly ignorant about possibilities for virtual machinery until computer systems engineering in the 20th Century introduced both new opportunities and new motivations for designing and building increasingly sophisticated types of virtual machinery. The majority of scientists and philosophers, and even many computer scientists, are still ignorant about what has been learnt and its scientific and philosophical (metaphysical) significance, partly summarised in:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>

One of the motivations for the Meta-Morphogenesis project is the conjecture that many hard unsolved problems in Artificial Intelligence, philosophy, neuroscience and psychology (including problems that have not yet been noticed?) require us to learn from the sort of evolutionary history discussed here, namely the history of construction kits and their products, especially increasingly

complex and sophisticated information processing machines, many of which are, or depend on, virtual machines. The fact that some construction kits can produce running virtual machinery supporting processes that include learning, perceiving, experiencing, wanting, disliking, inventing, supposing, asking, deciding, intending, pleasure, pain, and many more gives them not only scientific but also philosophical/metaphysical significance: they potentially provide answers to old philosophical questions, some of which have been recently rephrased in terms of a notion of grounding (interpreted as "metaphysical causation" by Wilson (2015)] -- a related paper on metaphysical ground of consciousness (and other aspects of life) in evolution is in preparation.)

Some previously unnoticed functions and mechanisms of minds and brains, including the virtual machinery they use, may be exposed by the investigation of origins and unobvious intermediate "layers" in biological information processing systems, based on older construction kits.

Showing how the FCK makes its derivatives possible, including all the processes and products of natural selection, is a challenge for science and philosophy. This is a long-term research programme with a good chance of being progressive in the sense of Imre Lakatos (1980), rather than degenerative.

Later, this may explain how to overcome serious current limitations of AI (artificial intelligence), robotics, neuroscience and psychology, as well as philosophy of mind and philosophy of science. In principle, it could draw attention to previously unnoticed features of the physical universe that are important aspects of the FCK.

Note:

My ideas have probably been influenced in more ways than I recognise by Margaret Boden, whose work has linked AI/Cognitive Science to Biology over several decades, notably in her *magnum opus* published in 2006 by OUP, *Mind As Machine: A history of Cognitive Science* (Vols 1-2).

[CONTENTS](#)

Philosophical Background: What is science? Beyond Popper and Lakatos

How is it possible for very varied forms of life to evolve from lifeless matter, including a mathematical species able to make the discoveries presented in Euclid's *Elements*?¹ Explaining evolution of mathematical insight is much harder than explaining evolution of consciousness! (Even insects must be conscious of aspects of their surroundings, used in control of motivation and actions).²)

¹ <http://www.gutenberg.org/ebooks/21076>

² I have argued elsewhere that the concept of "consciousness" labelled by a noun is problematic, in part because the adjectival forms are more basic than the noun, and the adjectival forms have types of context-sensitivity that can lead to truth-values that depend on context in complex ways. Some of the issues are summarised in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html>

An outline explanation of the evolutionary possibilities is based on construction kits produced by evolution, starting from a "fundamental" construction kit provided by physics and chemistry. New construction kits make new things (including yet more powerful construction kits) possible.

The need for science to include theories that explain how something is possible has not been widely acknowledged. Explaining how X is possible (e.g. how humans playing chess can produce a certain board configuration) need not provide a basis for *predicting* when X will be realised, so the theory used cannot be falsified by non-occurrence. [Popper, \[1934\]](#) labelled such theories "non-scientific" - at best metaphysics. His falsifiability criterion has been blindly followed by many scientists who ignore 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. Later, Popper shifted his ground, e.g. in [\[Popper 1978\]](#), and expressed great admiration for Darwin's theory of Natural Selection, despite its unfalsifiability.

[Lakatos \[1980\]](#) extended Popper's philosophy of science, showing how to evaluate competing scientific research programmes over time, according to their progress. He offered criteria for distinguishing "progressive" from "degenerating" research programmes, on the basis of their patterns of development, e.g. whether they systematically generate questions that lead to new empirical discoveries, and new applications. It is not clear to me whether he understood that his distinction could also be applied to theories explaining how something is possible.

[Chapter 2](#) of [\[Sloman 1978\]³](#) modified the ideas of Popper and Lakatos to accommodate scientific theories about what is *possible*, e.g. types of plant, types of animal, types of reproduction, types of consciousness, types of thinking, types of learning, types of communication, types of molecule, types of chemical interaction, and types of biological information processing.

³<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2>

The chapter presented criteria for evaluating theories of what is possible and how it is possible, including theories that straddle science and metaphysics. Insisting on sharp boundaries between science and metaphysics harms both. Each can be pursued with rigour and openness to specific kinds of criticism.

A separate paper on "Using construction kits to explain possibilities"⁴ includes a section entitled "Why allowing non-falsifiable theories doesn't make science soft and mushy", and discusses the general concept of "explaining possibilities", its importance in science, the criteria for evaluating such explanations, and how this notion conflicts with the falsifiability requirement for scientific theories. Further examples are in [\[Sloman 1996a\]](#). The extremely ambitious Turing-inspired Meta-Morphogenesis project⁵ first proposed in [\[Sloman 2013b\]](#), depends on these ideas, and will be a test of their fruitfulness, in a combination of metaphysics and science.

⁴ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/explaining-possibility.html>

⁵ Summarised in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

This paper, straddling science and metaphysics, asks: *How is it possible for natural selection, starting on a lifeless planet, to produce billions of enormously varied organisms, living in environments of many kinds, including mathematicians able to discover and prove geometrical and topological theorems?* I emphasise the biological basis of such mathematical abilities (a) because current Artificial Intelligence techniques do not seem to be capable of explaining and replicating them (and neither, as far as I can tell, can current neuroscientific theories) and (b) because those mathematical competences may appear to be esoteric peculiarities of a tiny subset of individuals, but arise out of deep features of animal and human cognition, that James Gibson came close to

noticing in his work on perception of affordances in [[Gibson 1979](#)]. Connections between perception of affordances and mathematical discoveries are illustrated in various documents on this web site. E.g. see [[Note 13](#)].

Explaining how discoveries in geometry and topology are possible is not merely a problem related to a small subset of humans, since the relevant capabilities seem to be closely related to mechanisms used by other intelligent species, e.g. squirrels and nest-building birds, and also in pre-verbal human toddlers, when they discover and use what I call "toddler theorems" [[Sloman 2013c](#)]. If this is correct, human mathematical abilities have an evolutionary history that precedes humans. Understanding that history may lead to a deeper understanding of later products.

A schematic answer to how natural selection can produce such diverse results is presented in terms of *construction kits*: the [Fundamental](#) (physical) Construction Kit (the FCK), and a variety of "concrete", "abstract" and "hybrid" [Derived](#) Construction Kits (DCKs). that together are conjectured to explain how evolution is possible, including evolution of mathematical abilities in humans and other animals, though many details are still unknown. The FCK and its relations to DCKs are crudely depicted below in Figures [FCK](#) and [DCK](#). Inspired by ideas in [[Kant 1781](#)], construction kits are also offered as providing Biological/Evolutionary foundations for core parts of mathematics, including mathematical structures used by evolution long before there were human mathematicians.

Note on "Making Possible": "X makes Y possible" as used here does not imply that if X does not exist then Y is impossible, only that *one* route to existence of Y is via X. Other things can also make Y possible, e.g., an alternative construction kit. So "makes possible" is a relation of sufficiency, not necessity. The exception is the case where X is the FCK - the *Fundamental Construction Kit* - since all concrete constructions must start from it (in this universe -- alternative "possible" universes based on different fundamental construction kits will not be considered here). If Y is abstract, there need not be something like the FCK from which it must be derived. The space of abstract construction kits may not have a fixed "root". However, the abstract construction kits that can be thought about by physically implemented thinkers may be constrained by a future replacement for the Church-Turing thesis, based on later versions of ideas presented here.

Although my questions about explaining possibilities arise in the overlap between philosophy and science [[Sloman 1978](#), Ch.2], I am not aware of any philosophical work that explicitly addresses the hypotheses presented and used here, though there seem to be examples of potential overlap, e.g. [[Bennett 2011](#), [Wilson 2015](#), [Dennett 1995](#)].

There are many overlaps between my work and the work of Daniel Dennett, as well as important differences. In particular, [[Dennett 1995](#)] comes close to addressing the problems discussed here, but as far as I can tell he nowhere discusses the need for evolved specific sorts of construction kit (fundamental and derived/evolved, concrete and abstract) to account for the continual availability of new options for natural selection) or the difficulty of replicating human and animal mathematical abilities (especially geometric and topological reasoning abilities) in (current) AI systems, and therefore does not mention the need to gain clues as to the nature of those abilities by investigating relevant evolutionary and developmental trajectories, as proposed here, because of hypothesised deep connections between those mathematical abilities and more general animal abilities to perceive and reason about affordances [[Gibson 1979](#)]. (I am grateful to Susan Stepney for reminding me of the overlap with Dennett after she read an earlier version of this paper. My impression is that he would regard all the different types of construction-kit discussed here as being

on the "cranes" side of his "cranes vs skyhooks" metaphor -- the other side, skyhooks, being intended to accommodate the sorts of divine intervention postulated by Dennett's opponents, with whom I also disagree.) I also think Dennett has not understood the full significance and causal powers of interactive portions of complex virtual machines, although he refers to them from time to time, often often in deprecating terms, using a spurious comparison between virtual machines and centres of gravity. But this is not the place to discuss his arguments: we are basically on the same side in several major debates.

A more detailed discussion of the idea of "making possible" is in a (draft -- Summer 2016) paper on how biological evolution made various forms of consciousness possible ("evolutionary grounding of consciousness"):

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/grounding-consciousness.html>

CONTENTS

1 A Brief History of Construction-kits

New introduction added 8 Apr 2016

Many thinkers, including biologists, physicists, and philosophers have attempted to explain how life as we know it could emerge on an initially lifeless planet. One widely accepted answer, usually attributed to Charles Darwin and Alfred Wallace is (in the words of Graham Bell(2008)):

Bell's answer:

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

But that does not explain what makes possible the options between which selections are made, apparently available in some parts of the universe (e.g. on our planet) but not others (e.g. the centre of the sun). Earlier versions of this paper suggested that evolutionary biologists have mostly failed to notice, or have ignored, this problem. The work of Kirschner [\[quoted below\]](#) seems to be an exception, though it is not clear that he recognised the requirements for layered construction kits discussed here.

Some thinkers have tried to explain how some of the features of physics and chemistry suffice to explain some crucial feature of life. E.g. [\[Schrödinger 1944\]](#) gave a remarkably deep and prescient analysis showing how features of quantum physics could explain how genetic information expressed in complex molecules composed of long aperiodic sequences of atoms could make reliable biological reproduction and development possible, despite all the physical forces and disturbances affecting living matter. Without those quantum mechanisms the preservation of precise complex genetic information across many generations during reproduction, and during many processes of cell division in a developing organism would be inexplicable, especially in view of the constant thermal buffeting by surrounding molecules. (So quantum mechanisms can help with local defeat of entropy?)

Portions of matter encoding complex biological "design" information need to be copied with great accuracy during reproduction in order to preserve specific features across generations. Design specifications also need to maintain their precise structure during all the turmoil of development, growth, etc. In both cases, small variations may be tolerated and may even be useful in allowing

minor tailoring of designs. (As explained in connection with parametrization of design specifications [below](#).) However if a change produces shortness on the left side of an animal without producing corresponding shortness on the right side, the animal's locomotion and other behaviours may be seriously degraded. The use of a common specification that controls the length of each side would make the design more robust, while avoiding the rigidity of a fixed length that rules out growth of an individual or variation of size in a species.

Schrödinger's discussion of requirements for encoding of biological information seems to have anticipated some of the ideas about requirements for information-bearers later published by Shannon, and probably influenced the work that led to the discovery of the double helix structure of DNA.

But merely specifying the features of matter that support forms of life based on matter does not answer the question: how do all the required structures come to be assembled in a working form? It might be thought that that could be achieved by chance encounters between smaller fragments of organisms producing larger fragments, and sometimes producing variants "preferred" by natural selection.

(Compare the demonstration of "Droguli" by Lionel and Roger Penrose mentioned in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html#new-droguli>)

But any engineer or architect knows that assembling complex structures requires more than just the components that appear in the finished product: all sorts of construction kits, tools, forms of scaffolding, and ancillary production processes are required. I'll summarise that by saying that biological evolution would not be possible without the use of construction kits and scaffolding mechanisms, some of which are not parts of the things they help to create. A more complete investigation would need to distinguish many sub-cases, including the kits that are and the kits that are not produced or shaped by organisms that use them.

The vast majority of organisms can neither assemble themselves simply by collecting inorganic matter or individual atoms or sub-atomic components from the environment, nor maintain themselves by doing so. Exceptions are organisms that manufacture all their nourishment from water and carbon-dioxide using photosynthesis, and the (much older) organisms that use various forms of chemosynthesis to produce useful organic molecules from inorganic molecules: [\[Bakermans \(2015\), Seckbach & Rampelotto \(2015\)\]](#) often found in extremely biologically hostile environments (hence the label "extremophiles"), some of which may approximate conditions before life existed.

Unlike extremophiles, most organisms, especially larger animals, need to acquire part-built fragments of matter of various sorts by consuming other organisms or their products. Proteins, fats, sugars, and vitamins are well known examples. This food is partly used as a source of accessible energy and partly as a source of important complex molecules needed for maintenance, repair, and other processes. Symbiosis provides more complex cases where each of the species involved may use the other in some way, e.g. as scaffolding or protection against predators, or as a construction mechanism, e.g. synthesising a particular kind of food harvested by the other species.

These examples show that unlike human-made construction kits, many biological construction kits are not created by their users but were previously produced by biological evolution without any influence from the users. This can change after an extended period of symbiosis or domestication, which can influence evolution of both providers and users. Compare the role of mate-selection

mechanisms based on previously evolved cognitive mechanisms that later proved useful for evaluating potential mates.

Even plants that use photosynthesis to assemble energy rich carbohydrate molecules from inorganic molecules (water and carbon dioxide) generally need to grow in an environment that already has organic waste matter produced by other living things, as any gardener, farmer, or forester knows.

All this implies that an answer to our main question must not merely explain how physical structures and processes suffice to make up all living organisms and explain all their competences. The physical world must also explain the possibility of the required "*construction kits*", including cases where some organisms are parts of the construction kits used by others. This paper attempts to explain, at a very high level of generality, some of what the construction kits do, how they evolve partly in parallel with the organisms that require them, and how importantly different types of construction kit are involved in different evolutionary and developmental processes, including construction kits concerned with information processing. This paper takes only a small number of small steps down a very long, mostly unexplored road.

NOTE 26 Jan 2017

I am grateful to Anthony Durity for drawing my attention to the work of Kirschner who clearly recognized the gap that was ignored in the quotation by [Graham Bell above](#). Some important steps towards providing answers seem to have been taken in [M.W. Kirschner and J.C. Gerhart \(2005\)](#). See also this interview with Marc Kirschner:

<http://bmcbiol.biomedcentral.com/articles/10.1186/1741-7007-11-110>

"So I think that to explain these developments in terms of the properties of cell and developmental systems will unify biology into a set of common principles that can be applied to different systems rather than a number of special cases that have to be learned somehow by rote." There is a useful short review here

<http://www.americanscientist.org/bookshelf/pub/have-we-solved-darwins-dilemma>

This seems to be consistent with the theory of evolved construction kits being developed as part of the M-M project, but does not answer all the questions, especially my driving question about how evolution produced mathematicians like Archimedes and Euclid.

A corollary: a tree/network of evolved construction kits

Added 10 Oct 2016

A more detailed discussion will need to show how construction kits of various sorts are located in a branching tree (or network) of kits produced and used by biological evolution and its products.

This will add another network to the previously discussed networks of evolved designs and evolved niches (sets of requirements) [\[Sloman 1995\]](#). A third network, produced in parallel with the first two will include evolved construction kits, including the fundamental construction kit and derived kits. These are unlikely to be simple tree-like networks because of the possibility of new designs, new niches and new construction kits being formed by combining previously evolved instances: so that branches merge.

[To be expanded.]

Beyond supervenience to richer forms of support

This suggested relationship between living things and physical matter goes beyond standard notions of how life (including animal minds) might *supervene* on physical matter, or how life might be *reduced* to matter.

According to the theory of construction kits, the existence of a living organism is not fully explained in terms of how all of its structures and capabilities are built on physical structures forming its parts, and their capabilities. For organisms also depend on the previous and concurrent existence of structures and mechanisms that are essential for the biological production processes. Of course, in many cases those are parts of the organism, and in other cases parts of the organism's mother. Various types of symbiosis add complex mutual dependencies to the story.

The above ideas are elaborated below in the form of a (still incomplete and immature) theory of construction kits, including construction kits produced by biological evolution and its products. I'll present some preliminary, incomplete, ideas about types of construction kit, and their roles in the astounding creativity of biological evolution. This includes the Fundamental Construction Kit (**FCK**) provided by the physical universe (including chemistry), and many Derived Construction Kits (**DCKs**) produced by physical circumstances, evolution, development, learning, ecosystems, cultures, etc.

We also need to distinguish construction kits of different types, such as:

- **concrete (physical) construction kits**, all of whose parts are made of physical materials, (including chemical structures and processes),
- **abstract construction kits** made of abstractions that can be instantiated in different ways e.g. grammars, axiom systems,
- **hybrid construction kits** with both concrete and abstract components (e.g. games like tennis, reasoning mechanisms).
- **meta-construction kits** that can be used to construct new construction kits.

Some construction kits include **scaffolding** subsystems, used during construction then later discarded.

The role(s) of information in life and its evolution

Life involves **information** in many roles, including reproductive information in a genome, internal information processing for controlling internal substates, and many forms of human information processing, including some that were not directly produced by evolution, but resulted from human learning or creation, possibly followed by teaching, learning or imitation, all of which require **information processing** capabilities. Such capabilities require concrete (physical) construction kits that can build and manipulate structures representing information: the combinations are hybrid construction kits. Understanding the types of construction kit that can produce information processing mechanisms of various types is essential for the understanding of evolution, especially evolution of minds of many kinds. In doing all this, evolution (blindly) made use of increasingly complex mathematical structures. Later humans and other animals also discovered and used some of the mathematical structures, without being aware of what they were doing. Later still, presumably after extensions to their information processing architectures supported additional metacognitive information processing, a subset of humans began to reflect on and argue about such discoveries and human mathematics was born (e.g. Euclid, and his predecessors). What they discovered included pre-existing mathematical structures, some of which they had used unwittingly

previously. This paper explores these ideas about construction kits, and some of their implications.

NOTES:

Examples of metacognitive mathematical or proto-mathematical information processing can be found in [\[Sloman 2015\]](#).

The word "information" is used here (and in other papers on this web site) not in the sense of Claude Shannon (1948), but in the much older *semantic* sense of the word, as used, for example, by Jane Austen. See: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html>

Added 12 Mar 2016

I have just acquired a copy of [\[Kauffman 1993\]](#). It seems to be directly relevant to this project, especially as regards the construction kits that build directly on chemical mechanisms.

[CONTENTS](#)

2 Fundamental and Derived Construction Kits (FCK, DCKs)

Natural selection alone cannot explain how evolution happens, for it must have options to select from. What sorts of mechanism can produce options that differ so much in so many ways, allowing evolution to produce microbes, fungi, oaks, elephants, octopuses, crows, new niches, ecosystems, cultures, etc.? Various sorts of construction kit, including evolved/derived construction kits, help to explain the emergence of new options.

What explains the possibility of these construction kits? Ultimately, features of fundamental physics, including those emphasised in [\[Schrödinger 1944\]](#) and discussed below (perhaps one of the deepest examples of a scientific theory attempting to explain how something is possible). Why did it take so much longer for evolution to produce baboons than bacteria? Not merely because baboons are more complex, but also because evolution had to produce more complex construction kits, to make baboon-building possible.

NOTE:

29 Dec 2016

One of the questions not discussed here is whether the use of construction kits merely allowed certain designs to be achieved faster, or whether the constructions were absolutely impossible without the kits. Compare this question: Without cranes, scaffolding, and other aids to skyscraper production would it have been possible for humans to create 21st century skyscrapers (perhaps as termites “grow” their cathedrals)? There seem to be several reasons why the answer must be negative, including, perhaps, the impossibility of creating very tall buildings resistant to very strong winds and minor earthquakes, and containing tons of equipment as well as humans, without using separately constructed parts, including girders that are part of the structure and temporary scaffolding required during construction. As far as I know, termites do not construct major portions of their cathedrals in separate places then bring them together for assembly.

What makes all of this possible is the construction kit provided by fundamental physics, the Fundamental Construction Kit (FCK) about which we still have much to learn, even if modern physics has got beyond the stage lampooned in this SMBC cartoon:

Enjoy the SMBC comic-strip comment on "fundamentality"

<http://www.smbc-comics.com/?id=3554>




Figure SMBC:Fundamental

Click the above to view the full 'comic strip',
or use this link to the image (and expand it in your browser):

<http://www.smbc-comics.com/comics/20141125.png>

(I am grateful to Tanya Goldhaber for the link.)

Perhaps SMBC will one day produce a similar cartoon whose dialogue ends thus:

Student: "Professor, what's an intelligent machine?"

Professor: "Anything smarter than what was intelligent a generation ago."

As hinted by the cartoon, there is not yet agreement among physicists as to what exactly the FCK is, or what it can do. Perhaps important new insights into properties of the FCK will be among the long term outcomes of our attempts to show how the FCK can support all the DCKs required for developments across billions of years, and across no-one knows how many layers of complexity, to produce animals as intelligent as elephants, crows, squirrels, or even humans (or their successors).

Construction-kits are the "hidden heroes" of evolution. Life as we know it requires construction kits supporting construction of machines with many capabilities, including growing many types of material, many types of mechanism, many different highly functional bodies, immune systems, digestive systems, repair mechanisms, reproductive machinery, information processing machinery (including digital, analogue and virtual machinery) and even producing mathematicians, such as Euclid and his predecessors!

A kit needs more than basic materials. If all the atoms required for making a loaf of bread could somehow be put into a container, no loaf could emerge. Not even the best bread making machine, with paddle and heater, could produce bread from atoms, since that would require atoms pre-assembled into the right amounts of flour, sugar, yeast, water, etc. Only different, separate,

histories can produce the molecules and multi-molecule components, e.g. grains of yeast or flour.


Likewise, no fish, reptile, bird, or mammal could be created simply by bringing together enough atoms of all the required sorts; and no machine, not even an intelligent human designer, could assemble a functioning airliner, computer, or skyscraper directly from the required atoms. Why not, and what are the alternatives? We first state the problem of constructing very complex working machines in very general terms and indicate some of the variety of strategies produced by evolution, followed later by conjectured features of a very complex, but still incomplete, explanatory story.

2.1 Combinatorics of construction processes

Reliable construction of a living entity requires: appropriate types of matter, machines that manipulate matter, physical assembly processes, stores of energy for use during construction, and usually *information* e.g. about which components to assemble at each stage, how to assemble them, and how to decide in what order to do so. This requires, at every stage, at least: (i) components available for the remaining stages, (ii) information about which components, and which movements and other processes are required, (iii) mechanisms capable of assembling the components, (iv) mechanisms able to decide what should happen next, whether selection is random or information-based. Sometimes the mechanisms are part of the structure assembled, sometimes not. Some may be re-usable multi-purpose mechanisms, while others are temporary structures discarded after use, e.g. types of scaffolding, and other tools. In this context the egg-shell of a bird or reptile could be viewed as scaffolding (a temporary structure supporting safe development).


The mechanisms involved in construction of an organism can be thought of as a construction kit, or collection of construction kits. Some components of the kit are parts of the organism and are used throughout the life of the mechanism, e.g. cell-assembly mechanisms used for growth and repair. Construction kits used for building information processing mechanisms may continue being used and extended long after birth as discussed in the section on epigenesis [below](#). All of the construction kits must ultimately come from the Fundamental Construction kit (FCK) provided by physics and chemistry.

Figure 1 FCK: The Fundamental Construction Kit



A crude representation of the Fundamental Construction Kit (FCK) (on left) and (on right) a collection of trajectories from the FCK through the space of possible trajectories to increasingly complex mechanisms.

Figure 2 DCK: Derived Construction Kits



Further transitions: a fundamental construction kit (FCK) on left gives rise to new evolved "derived" construction kits, such as the DCK on the right, from which new trajectories can begin, rapidly producing new more complex designs, e.g. organisms with new morphologies and new information processing mechanisms. The shapes and colours (crudely) indicate qualitative differences between components of old and new construction kits, and related trajectories. A DCK trajectory uses larger components and is therefore much shorter than the equivalent FCK trajectory.

The Fundamental Construction Kit (FCK) provided by the physical universe made possible all the forms of life that have so far evolved on earth, and also possible, but still unrealised, forms of life, in possible types of physical environment. Fig. 1 shows (crudely) how a common initial construction kit can generate many possible trajectories, in which components of the kit are assembled to produce new instances (living or non-living).

The space of possible trajectories for combining basic constituents is enormous, but routes can be shortened and search spaces shrunk by building derived construction kits (DCKs), that are able to assemble larger structures in fewer steps⁶, as indicated in Fig. 2.

⁶ Assembly mechanisms are part of the organism, illustrated in a video of grass "growing itself" from seed

<https://www.youtube.com/watch?v=JbiQtfr6AYk>.

In mammals with a placenta, more of the assembly process is shared between mother and offspring.

The history of technology, science, engineering and mathematics includes many transitions in which new construction kits are derived from old ones by humans. That includes the science and technology of digital computation, where new advances used an enormous variety of discoveries and inventions including these (among many others):

1. Jacquard looms in which punched cards were used to control operations in complex weaving machines.
2. punched cards, punched tape, and mechanical sorting devices in business data-processing;

3. electronic circuits, switches, mercury delay lines, vacuum tubes, switchable magnets, and other devices;
4. arrays of transistors, connected electronically;
5. machine language instructions expressed as bit-patterns, initially laboriously "loaded" into electronic computers by making connections between parts of re-configurable circuits, and, in later systems, by setting banks of switches on or off;
6. 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;
7. compilers and assemblers that translate symbolic programs into bit patterns;
8. use of operating systems: including programs that manage other programs and hardware resources;
9. electronic interfaces between internal digital mechanisms and external input, output, and networking devices, along with operating system extensions and specialised software interface-handlers to deal with interactions between internal and external events and processes;
10. many types of higher level programming language that are compiled to machine language or to intermediate level languages before programs start running;
11. 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.
12. Network-level management systems allowing ever more complex collections of computers and other machines to be connected to support many kinds of control, collaboration, and communication.

Abstraction via parametrization

Particular inventions were often generalised and parametrized, using mathematical abstractions, by replacing parts of a specification for a design with slots, or gaps, or "variables", that could have different "fillers" or "values". This produces patterns or abstract designs, that can be reused in new contexts by providing different specifications to fill the slots in different designs.

This is a form of discovery with a long history in mathematics, e.g. discovering a "Group" pattern in different areas of mathematics and then applying general theorems about groups to particular instantiations.

During the growth of an organism an unchanging overall design needs to accommodate systematically changing physical parts and corresponding systematically changing control features that make good use of increased size, weight, strength, reach, and speed, without, for example, losing control because larger parts have greater momentum.

Note (added 9 Apr 2016):

According to Wikipedia, a new born foal "will stand up and nurse within the first hour after it is born, can trot and canter within hours, and most can gallop by the next day". Contrast hunting mammals and humans, which are born less well developed and go through more varied forms of locomotion, requiring more varied forms of control and learning, during relative and absolute changes of sizes of body parts, relations between sizes, strength of muscles, perceptual competences, and predictive capabilities.

<https://en.wikipedia.org/wiki/Foal>

Control mechanisms for a form of movement, such as crawling, walking, or running, will continually need to adapt to changes of various features in physical components (size, bone strength, muscle strength, weight, joint configuration, etc.) It is unlikely that every such physical change leads to a complete revision of control mechanisms: it is more likely that for such species evolution produced parametrized control so that parameters can change while the overall form of control persists for each form of locomotion.

If such mathematically abstract control structures have evolved within a species they could perhaps also have been re-used across different genomes for derived species that vary in size, shape, etc. but share the abstract topological design of a vertebrate. Similar comments about the need for mathematical abstractions for control apply to many actions requiring informed control, e.g. grasping, pushing, pulling, breaking, feeding, carrying offspring, etc.

Related ideas are mentioned in connection with construction kits for ontologies [below](#).

The discovery by evolution of designs that can be parametrized and re-used in different contexts involves discovery and use of mathematical structures that can be instantiated in different ways. This is one of the main reasons for regarding evolution as a "Blind Theorem Prover". Evolution implicitly proves many such forms of mathematical generalization to be possible and useful. The historical evolutionary/developmental trajectory leading up to a particular instantiation of such a possibility constitutes an implicit **proof** of that possibility. Of course all that can happen without any explicit recognition that it has happened: that's why evolution can be described as a "Blind mathematician" in the same sense of "Blind" as has been used in the label "Blind watchmaker", e.g. by Richard Dawkins (echoing Paley). https://en.wikipedia.org/wiki/The_Blind_Watchmaker.

Note added 15 Dec 2018

There are several different theories and computational models that make use of the fact that when searching a space of designs, or solutions to problems, instead of investigating only different combinations of the smallest building blocks (e.g. bits in the case of a bit string), it is often much more efficient to build new search spaces from larger building blocks that have already been found to be useful. An example use of this idea is the theory of Genetic Programming.

New tools, including virtual machinery

The production of new applications also frequently involved production of new tools for building more complex applications, and new kinds of virtual machinery, allowing problems, solutions and mechanisms to be specified in a manner that was independent of the physical implementation details.

Natural selection did something similar on an even larger scale, with far more variety, probably discovering many obscure problems and solutions (including powerful abstraction) still unknown to us. (An educational moral: teaching only what has been found most useful can discard future routes to possible major new advances - like depleting a gene pool.)

Biological construction kits derived from the FCK can combine to form new Derived Construction Kits (DCKs), some specified in genomes, and (very much later) some discovered or designed by individuals (e.g. during epigenesis Sect. [2.3](#)), or by groups, for example new languages. Compared with derivation from the FCK, the rough calculations above show how DCKs can enormously speed up searching for new complex entities with new properties and behaviours. See Fig. [2](#).

New DCKs that evolve in different species in different locations may have overlapping functionality, based on different mechanisms: a form of *convergent evolution*. E.g., mechanisms enabling elephants to learn to use trunk, eyes, and brain to manipulate food may share features with those enabling primates to learn to use hands, eyes, and brains to manipulate food. In both cases, competences evolve in response to structurally similar affordances in the environment. This extends ideas in [Gibson 1979] to include affordances for a species, or collection of species.⁷

⁷ Implications for evolution of vision and language are discussed in <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>

NOTE on vision (Added 3 Sep 2017)

There may also be closely related, partially overlapping, affordances, including for example the affordances suited to compound eyes, with multiple lenses each producing small amounts of visual information, which evolved a long time before simple eyes with a single lens projecting a rich spatially organised information structure onto a retina with multiple receptors with different sized receptive fields. Both types are used with extraordinary success in the species that have them, but the requirements and opportunities for information processing are very different.

This implies that evolution had to produce not only different construction kits for producing and assembling the **physical** components of the two kinds of eyes, but also corresponding construction kits for producing the two kinds of **information processing system**.

Moreover, each of the two kinds evolved many different specialisations. E.g. some mammals and birds have eyes with overlapping receptive fields providing opportunities for using triangulation to infer distance (stereopsis), whereas others have non-overlapping (or barely overlapping) visual fields, providing information about more of the surrounding environment.

Readers are invited to think about how the differences are relevant to different needs of hunters and prey, as well as differing needs of nest/shelter builders and animals that don't build or use nests.

Combinatorics of biological constructions.

If there are N types of basic component and a task requires an object of type O composed of K basic components, the size of a blind exhaustive search for a sequence of types of basic component to assemble an O is up to N^K sequences, a number that rapidly grows astronomically large as K increases. If, instead of starting from the N types of *basic* component, the construction uses M types of *pre-assembled* component, each containing P basic components, then an O will require only K/P pre-assembled parts. The search space for a route to O is reduced in size to $M^{(K/P)}$.

Compare assembling an essay of length 10,000 characters (a) by systematically trying elements of a set of about 30 possible characters (including punctuation and spaces) with (b) choosing from a set of 1000 useful words and phrases, of average length 50 characters. In the first case each choice has 30 options but 10,000 choices are required. In the second case there are 1000 options per choice, but far fewer stages: 200 instead of 10,000 stages. So the size of the (exhaustive) search space is reduced from 30^{10000} , a number with 14,773 digits, to about 1000^{200} , a number with only 602 digits: a very much smaller number. So trying only good pre-built substructures at each stage of a construction process can produce a huge reduction of the search space for solutions of a given size - though some solutions may be missed. If no useful result is found the

search has to go back to "starting from fundamentals", which may take a very long time, for the reasons given.

So, learning from experience by storing useful subsequences can achieve dramatic reductions, analogous to a house designer moving from thinking about how to assemble atoms, to thinking about assembling molecules, then bricks, planks, tiles, then pre-manufactured house sections.

The reduced search space contains fewer samples from the original possibilities, but the original space is likely to have a much larger proportion of useless options. As sizes of pre-designed components increase, so does the variety of pre-designed options to choose from at each step, though far, far, fewer search steps are required for a working solution, if one exists: a very much shorter evolutionary process.

The cost of searching in the the reduced space may be exclusion of some design options. In the case of biological evolution, there is not one process of design and construction: huge numbers of evolutionary trajectories are explored at different times and locations, and occasionally natural disasters that wipe out a collection of very successful solutions may force a return to a different branch in the search space that leads to even greater successes, e.g. producing more versatile or more intelligent species, as seems to have happened after the destruction of most dinosaurs. (CHECK)

This indicates intuitively, but very crudely, how using increasingly large, already tested useful part-solutions can enormously reduce the search for viable solutions -- if they exist!

Comparison with use of Memo-functions in computing

The technique of storing previously found results of computations is familiar to many programmers, for example in the use of "memo-functions" ("memoization") to reduce computation time. As far as I know the idea was first published in [Michie \(1968\)](#), building on work by Robin Popplestone.

Example

The Fibonacci sequence is defined as a sequence of whole numbers starting with 0 and 1, and thereafter adding the last two computed numbers to get the next number. I.e. the function fib is defined as follows

```
Definition of "fib(X)"
fib(0) = 0
fib(1) = 1
if X > 1 then fib(X) = fib(X-2) + fib(X-1)
```

From that you can work out that $\text{fib}(2) = \text{fib}(0) + \text{fib}(1)$, i.e. 1.

However, the number of steps required to compute the output number for each input number keeps getting bigger and bigger much faster than the input numbers do. (Try using that definition to compute $\text{fib}(4)$, $\text{fib}(5)$ and $\text{fib}(6)$, giving results 3, 5, and 8, respectively, but requiring rapidly increasing numbers of computational steps.)

Readers used to programming with extendable arrays can work out how to use an array that initially contains no values, then, as each invocation of fib with a new number, e.g. N, produces a new value e.g. V, the value V can be stored as the N'th item of the array. That will make it possible to redefine fib so that after it is asked to calculate the value for any number, e.g. 20, when the largest previous value calculated was for 13, it will very quickly calculate the values

for 14, 15, .. 19 and then use the last two to calculate the value for 20. The result is achieved in a tiny fraction of the time required without the extra storage.

This technique is called use of "Memo-functions", Readers who are unfamiliar with this should try using the above definition to compute fib(2), fib(3), fib(4), to get a feel for how the computation required for each new number expands, and also how much repetition is involved.

to use everything previously constructed, as in the Fibonacci sequence, where the Nth number is the sum of previous two Fibonacci numbers, starting with fib(0) = 0 and fib(1) = 1. So computing fib(100) requires computing the two previous values each of which requires its two previous values, and so on. The time required grows surprisingly fast, with very many repeated computations. Programmers have learnt to produce programs that avoid re-computing, and instead simply remember computed values. Then to get the Nth number after computing the previous two numbers all one has to do is add the last two stored values. If they have not yet been computed the mechanism simply descends to the largest stored number and computes the missing intermediate ones, which can produce an enormous saving. These are called "memo-functions", because they use a memory of previous computations.

Structure-sharing and Genetic Programming

Use of memo-functions is one of a class of strategies that can produce significant savings in searching or computing, by storing potentially re-usable partial solutions to sub-problems while trying to solve a complex problem. If the re-usable techniques or designs can be indexed in some way that allows opportunities for re-using them to be identified, this may allow later sub-problems to be solved far more quickly than if all partial solutions have to be re-discovered "bottom-up" in order to be re-used. In biological evolution this could include making part of the genome that was originally found to be useful in certain contexts available to be tried out in new contexts that partially match the old contexts.

The extent to which this is useful will depend on the extent to which the universe (or some inhabited part of it) is "cognitively friendly", in the terminology of [Sloman\(1978, Chap.9\)](#). In a cognitively friendly environment, trying accommodate old solutions to new problems is often far more efficient than always searching "bottom-up" (try simplest new steps first) for new solutions. This will be true of an environment in which certain classes of structure larger than atoms or sub-atomic particles tend to recur, as is common in products of biological evolution. (So biological evolution tends to change a portion of the universe in ways that increase the opportunities for evolution.)

This re-use technique, sometimes referred to as "structure sharing", is often re-invented by AI researchers. The use of "crossover" in sexual reproduction (and in Genetic Algorithms), is partly analogous insofar as it allows parts of each parent's design specification to be used in new combinations. The family of computational search techniques known as "Genetic Programming"⁸ uses related ideas. (I suspect evolution has far more re-use mechanisms than have so far been identified.)

The use of various sorts of construction kits discussed in this paper can also be seen as a variety of second-order structure sharing, insofar as a construction kit serves as a solution (or family of related solution components) usable in different contexts, e.g. in different members of a species, or in members of related species, or at particular stages during development. But the construction kit is used as an intermediary to produce the mechanisms actually used by individual organisms to

meet biological needs or solve problems.

⁸ https://en.wikipedia.org/wiki/Genetic_programming

Storing solutions vs storing information about solutions

In biological evolution, instead of previous *solutions* being stored for future re-use (as in a Museum), *information about how to build components of previous solutions* is stored in genomes. (This statement is not completely accurate, since the *interpretation* process in constructing a new individual also uses information implicit in the construction mechanisms and in the environment, as explained below in Section [2.3](#), on epigenesis and in [Figure EVO-DEVO](#)).

If appropriate construction kits and materials are available, information structures specifying designs (e.g. blue-prints for complex machines) are much cheaper and quicker to store, copy, modify and re-use than physical instances of those designs (the machines themselves).

Evolution, the Great Blind Mathematician, discussed further [below](#), seems to have discovered memoization and the usefulness of exploration in a space of *specifications* for working systems rather than the systems themselves, long before we did.

"Negative memoization" A closely related strategy is to record fragments that *cannot* be useful in certain types of problem, in order to prevent wasteful attempts to use such fragments. Expert mathematicians learn from experience which options are useless (e.g. dividing by zero). This could be described as "negative-memoization". Are innate aversions examples of evolution doing something like that? In principle it might be possible for a genome to include warnings about not attempting certain combinations of design fragments, but I know of no evidence that that happens. Possible mechanisms might include components of construction kits that detect certain patterns in a fertilized egg, and abort development of eggs that produce those patterns. (Humans have discovered a few techniques for doing this. The ethical problems of human intervention don't arise for mechanisms produced by evolution that are beyond human control.)

When prior information about useful components and combinations of pre-built components is not available, random assembly processes can be used. If mechanisms are available for recording larger structures that have been found to be useful or useless, the search space for new designs can be shrunk. By doing the searching and experimentation using *information* about how to build things rather than directly recombining the built physical structures themselves, evolution reduces the problem of *recording* what has been learnt.

Effects on genome size

NOTE: making faster or easier vs making possible:

29 Dec 2016

One of the questions not discussed in detail here is whether the use of construction kits merely allowed certain designs to be achieved faster, or whether certain sorts of construction are absolutely impossible without the use of kits containing structures that are not part of the final product.

Compare this question: Without cranes, scaffolding, and other aids to skyscraper production would it have been possible for humans to create 21st century skyscrapers (perhaps as termites "grow" their cathedrals)? There seem to be several reasons why the answer must be

negative, including, perhaps, the impossibility of creating very tall buildings resistant to very strong winds and minor earthquakes, and containing tons of equipment as well as humans, without using separately constructed parts, including girders that are part of the structure and temporary scaffolding required during construction.

As far as I know, termites do not construct major portions of their cathedrals in separate places then bring them together for assembly. However, many animals do use separately grown structures that happen to be around, e.g. fetching already available twigs, leaves, mud, stones, and hair of other animals, to build their nests. Do any go further and assemble separate structures to be used during nest-construction then discarded afterwards?

2.2 Construction Kit Ontologies

A construction kit (and its products) can exist without being described. However, scientists need to use various forms of language in order to describe the entities they observe or postulate in explanations, and to formulate new questions to be answered. So a physicist studying the FCK will need one or more construction kits for defining concepts, formulating questions, formulating theories and conjectures, constructing models, etc. Part of the process of science is extending construction kits for theory formation. Something similar must be done by natural selection: extending useful genetic information structures that store specifications for useful components.

This relates to claims that have been made about requirements for control systems and for scientific theories. For example, if a system is to be capable of distinguishing N different situations and responding differently to them, it must be capable of being in at least N different states (recognition+control states). This is a variant of Ashby's "Law of Requisite Variety" [Ashby 1956]. Many thinkers have discussed representational requirements for scientific theories, or for specifications of designs. [Chomsky 1965] distinguished requirements for theories of language, which he labelled *observational adequacy* (covering the variety of observed uses of a particular language), *descriptive adequacy* (covering the intuitively understood principles that account for the scope of a particular language) and *explanatory adequacy* (providing a basis for explaining how any language can be acquired on the basis of data available to the learner). These labels were vaguely echoed in [McCarthy Hayes 1969] who described a form of representation as being *metaphysically adequate* if it can express anything that can be the case, *epistemologically adequate* if it can express anything that can be known by humans and future robots, and *heuristically adequate* if it supports efficient modes of reasoning and problem-solving. (I have simplified all these proposals.)

Construction-kit powers.

Requirements can also be specified for powers of various sorts of biological construction kit. The fundamental construction kit (FCK) must have the power to make any form of life that ever existed or will exist possible, using huge search spaces if necessary. DCKs may meet different requirements, e.g. each supporting fewer types of life form, but enabling those life forms to be "discovered" in a shorter time by natural selection, and reproduced (relatively) rapidly. Early DCKs may support the simplest organisms that reproduce by making copies of themselves [Ganti 2003], or precursors of life that produce some the mechanisms required [Froese et al(2014)].

At later stages of evolution, DCKs are needed that allow construction of organisms that change their properties during development and change their control mechanisms appropriately as they grow [Thompson 1917]. This requires abilities to produce individuals whose features are

parametrised with parameters that change over time.

More sophisticated DCKs must be able to produce species whose members use epigenetic mechanisms to modify their knowledge and their behaviours not merely as required to accommodate their own growth but also to cope with changing physical environments, new predators, new prey and new shared knowledge. A special case of this is having genetic mechanisms able to support development of a wide enough range of linguistic competences to match any type of human language, developed in any social or geographical context. However, the phenomenon is far more general than language development, as discussed in the next section.

2.3 Construction kits built during development (epigenesis)

Some new construction kits are products of evolution of a species and are initially shared only between a few members of the species (barring genetic abnormalities), alongside cross-species construction kits shared between species, such as those used in mechanisms of reproduction and growth in related species. Evolution also discovered the benefits of "meta-construction-kits": mechanisms that allow members of a species to build new construction kits during their own development.

Examples include mechanisms for learning that are initially generic mechanisms shared across individuals, and 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, part way through learning 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 Sloman 2007](#)], summarised in [Figure EVO-DEVO](#). The meta-configured competences are partly specified in the genome but those partial, abstract, specifications are instantiated in combination with information abstracted from individual experiences in various domains, of increasing abstraction and increasing complexity.

Meta-Configured genomes

Mathematical development and language development in humans both seem to be special cases of growth of such meta-configured competences. Related ideas are in [[Karmiloff-Smith 1992](#)].

Multiple routes from genome to behaviours




Figure 3: Figure EVO-DEVO:

A particular collection of construction kits specified in a genome can give rise to very different individuals in different contexts if the genome interacts with the environment in increasingly complex ways during development, allowing enormously varied developmental trajectories. Precocial species use only the downward routes on the left, producing only "preconfigured" competences. Competences of members of "altricial" species, using staggered development, may be far more varied within a species. Results of using earlier competences interact with the genome, producing "meta-configured" competences shown on the right. This is a modified version of a figure in [Chappell Sloman 2007].

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. Nobody knows the full variety of ways in which parallel construction processes can exercise mutual control in developing organisms. One implication of [Figure EVO-DEVO](#) is that there are not always simple correlations between genes and organism features.

The main idea could be summarised approximately as follows:

Instead of the genome determining how the organism reacts to its environment, the environment can cumulatively determine how the genome expresses itself: with different sorts of influence at different stages of development. This should not be confused with theories that attempt to measure **percentages** of genetic vs environmental influence in individual development. Numerical measures in this context are much shallower than specifications of structures and their interactions. Compare: expressing the *percentage* of one composer's influence on another (e.g. Haydn's influence on Beethoven) would give little understanding of

what the later composer had learnt from his or her predecessor. Often emphasising measurement over precise description can obfuscate science instead of deepening it. Likewise emphasising *correlations* can get in the way of understanding *mechanisms*.

Explaining the many ways in which a genome can orchestrate parallel processes of growth, development, formation of connections, etc. is a huge challenge. A framework allowing abstract specifications in a genome to interact with details of the environment in instantiating complex designs is illustrated schematically in Fig. 3. An example might be the proposal in [Popper 1976] that newly evolved desires of individual organisms (e.g. desires to reach fruit in taller trees) could indirectly and gradually, across generations, influence selection of physical characteristics (e.g. longer necks, abilities to jump higher) that improve success-rates of actions triggered by those desires. Various kinds of creativity, including mathematical creativity, might result from such transitions. This generalises Waddington's "epigenetic landscape" metaphor [Waddington 1957], by allowing individual members of a species to partially construct and repeatedly modify their own epigenetic landscapes instead of merely following paths in a landscape that is common to the species. Mechanisms that increase developmental variability may also make new developmental defects possible (e.g. autism?)⁹.

⁹ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/autism.html>

2.4 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, they require increasingly sophisticated control mechanisms. New sorts of control often use new types of information. Processing that information may require new mechanisms. That may require new construction kits for building new types of information processing mechanism.

The simplest organisms use only a few types of (mainly chemical) sensor, providing information about internal states and the immediate external physical environment. They have very few behavioural options. They acquire, use and replace fragments of information, using the same forms of information throughout their life, to control deployment of a fixed repertoire of capabilities.

More complex organisms acquire information about enduring spatial locations in extended terrain, including static and changing routes between static and changing resources and dangers. They need to construct and use far more complex (internal or external) information stores about their environment, and, in some cases, "meta-semantic" information about information processing, in themselves and in others, e.g. conspecifics, predators and prey. (Some of the types of evolutionary transition towards increasing complexity of form and behaviour are shown schematically later in [Figure 4](#).)

What forms can all the intermediate types of information take? Many controlled systems have states that can be represented by a fixed set of physical measures, often referred to as "variables", representing states of sensors, output signals, and internal states of various sorts. Such systems have many engineering applications, so many researchers are tempted to postulate them in biological information processing. Are they adequate?

Relationships between static and changing state-components in such systems are often represented mathematically by equations, including differential equations, and constraints (e.g. inequalities) specifying restricted, possibly time-varying, ranges of values for the variables, or magnitude relations between the variables. A system with N variables (including derivatives) has a state of a fixed dimension, N. The only way to record new information in such systems is through static or changing values for numeric variables -- changing "state vectors", and possibly alterations in the equations.

A typical example of such an approach is the work of [Powers 1973], inspired by [Wiener 1961] and [Ashby 1952]. There are many well understood special cases, such as simple forms of homeostatic control using negative feedback. Neural net based controllers often use large numbers of variables clustered into strongly interacting sub-groups, groups of groups, etc. Are these structures and mechanisms adequate for all biological information processing -- including human perception and reasoning?

For many structures and processes, a set of numerical values and rates of change linked by equations (including differential equations) expressing their changing relationships is an adequate form of representation, but not for all, as implied by the discussion of types of *adequacy* in Section 2.2. That's why chemists use *structural* formulae, e.g. diagrams showing different sorts of bond between atoms, and collections of diagrams showing how bonds change in chemical reactions. Linguists, programmers, computer scientists, architects, structural engineers, map-makers, map-users, mathematicians studying geometry and topology, composers, and many others, work in domains where structural diagrams, logical expressions, grammars, programming languages, plan formalisms, and other *non-numerical* notations express information about structures and processes that is not usefully expressed in terms of collections of numbers and equations linking numbers.¹⁰

¹⁰ Examples include:

https://en.wikipedia.org/wiki/Parse_tree

https://en.wikipedia.org/wiki/Structural_formula

<https://en.wikipedia.org/wiki/Flowchart>

https://en.wikipedia.org/wiki/Euclidean_geometry

https://en.wikipedia.org/wiki/Entity-relationship_model

https://en.wikipedia.org/wiki/Programming_language

Of course, any information that can be expressed in 2-D written or printed notation, such as grammatical rules, parse trees, logical proofs, and computer programs, can also be converted into a large array of numbers by taking a photograph and digitising it. Although such processes are useful for storing or transmitting documents, they add so much irrelevant numerical detail that the original functions, such as use in checking whether an inference is valid, or manipulating a grammatical structure by transforming an active sentence to a passive one, or determining whether two sentences have the same grammatical subject, or removing a bug from a program, or checking whether a geometric construction proves a theorem, become inaccessible until the original non-numerical structures are extracted, often at high cost.

Similarly, collections of numerical values will not always adequately represent information that is biologically useful for animal decision making, problem solving, motive formation, learning, etc. Moreover, biological sensors are poor at acquiring or representing very precise information, and neural states often lack reliability and stability. (Such flaws can be partly compensated for by using

many neurons per numerical value and averaging.) More importantly, the biological functions, e.g. of visual systems, may have little use for absolute measures if their functions are based on *relational* information, such as that A is closer to B than to C, A is biting B, A is keeping B and C apart, A can fit through the gap between B and C, the joint between A and B is non-rigid, A cannot enter B unless it is reoriented, and many more.

As [Schrödinger 1944] pointed out, topological structures of molecules can reliably encode a wide variety of types of genetic information, and may also turn out to be useful for recording other forms of structural information. Do brains employ them? There are problems about how such information can be acquired, derived, stored, used, etc. [Chomsky 1965] pointed out that using inappropriate structures in models may divert attention from important biological phenomena that need to be explained-see Sect. 2.2, above.

Max Clowes, who introduced me to AI in 1969, made similar points about research in vision around that time.¹¹ So subtasks for this project include identifying biologically important types of non-numerical (e.g. relational) information content and ways in which such information can be stored, transmitted, manipulated, and used. We also need to explain how mechanisms performing such tasks can be built from the FCK, using appropriate DCKs.

¹¹ <http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#71>

NOTE: Refactoring of designs

Note added 9 Oct 2016, to be expanded

At a panel discussion in Edinburgh on "Soluble versus Insoluble : The Legacy of Alan Turing" 8th September 2016 <http://conferences.inf.ed.ac.uk/blc/> Alan Bundy pointed out that a feature of software development is "refactoring", described on Wikipedia as *"The process of restructuring existing computer code-changing the factoring-without changing its external behavior. Refactoring improves nonfunctional attributes of the software. Advantages include improved code readability and reduced complexity; these can improve source-code maintainability and create a more expressive internal architecture or object model to improve extensibility."*

https://en.wikipedia.org/wiki/Code_refactoring.)

That observation drew my attention to the need to expand this document to show that evolution's design changes also include refactoring, e.g. when a working design is parametrised to make it re-usable with different details. I'll add more about that here later.

Note:

The so-called "nonfunctional" attributes of designs are more accurately labelled "meta-functional" attributes as proposed in [Sloman & Vernon, 2007]

2.5 Increasingly varied mathematical structures

Electronic computers made many new forms of control possible, including use of logic, linguistic formalisms, planning, learning, problem solving, vision, theorem proving, teaching, map-making, automated circuit design, program verification, and many more. The world wide web is 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 types of control mechanism in computer-based

systems¹² now extend far beyond the numerical sorts familiar to control engineers.

¹² Often misleadingly labelled "non-linear".
http://en.wikipedia.org/wiki/Control_theory
http://en.wikipedia.org/wiki/Nonlinear_control

Organisms also need multiple control systems, not all numerical. A partially constructed percept, thought, question, plan or terrain description has parts and relationships, to which new components and relationships can be added and others removed, as construction proceeds and errors are corrected, building structures with changing complexity - unlike a fixed-size collection of variables assigned changing values. Non-numerical types of mathematics are needed for describing or explaining such systems, including topology, geometry, graph theory, set theory, logic, formal grammars, and theory of computation. A full understanding of mechanisms and processes of evolution and development may need new branches of mathematics, including mathematics of non-numerical structural processes, such as chemical change, or changing "grammars" for internal records of complex structured information. The importance of non-numerical information structures has been understood by many mathematicians, logicians, linguists, computer scientists and engineers, but many scientists still focus only on numerical structures and processes. They sometimes attempt to remedy their failures by using statistical methods, which, in restricted contexts, can be spectacularly successful, as shown by recent AI successes, whose limitations I have criticised elsewhere. ¹³

¹³ E.g. See [[Sloman 2015](#)]

The FCK need not be able to produce all biological structures and processes *directly*, in situations without life, but it must be rich enough to support successive generations of increasingly powerful DCKs that together suffice to generate all possible biological organisms evolved so far, and their behavioural and information processing abilities. Moreover, the FCK, or DCKs derived from it, must include abilities to acquire, manipulate, store, and use information structures in DCKs that can build increasingly complex machines that encode information, including non-numerical information. Since the 1950s we have also increasingly discovered the need for new *virtual* machines as well as *physical* machines [[Sloman 2010](#),[Sloman 2013a](#)].

Large scale physical processes usually involve a great deal of variability and unpredictability (e.g. weather patterns), and sub-microscopic indeterminacy is a key feature of quantum physics; yet, as [[Schrödinger 1944](#)] observed, life depends on genetic material in the form of very complex objects built from very large numbers of small-scale structures (molecules) that can preserve their *precise* chemical structure, over long time periods, despite continual thermal buffeting and other disturbances.

Unlike non-living natural structures, important molecules involved in reproduction and other biological functions are copied repeatedly, predictably transformed with great precision, and used to create very large numbers of new molecules required for life, with great, but not absolute, precision. This is *non-statistical* structure preservation, which would have been incomprehensible without quantum mechanics, as explained by Schrödinger.

That feature of the FCK resembles "structure-constraining" properties of construction kits such as Meccano, TinkerToy and Lego¹⁴ that support structures with more or less complex, discretely varied topologies, or kits built from digital electronic components, that also provide extremely

reliable preservation and transformations of precise structures, in contrast with sand, water, mud, treacle, plasticine, and similar materials. Fortunate children learn how structure-based kits differ from more or less amorphous construction kits that produce relatively flexible or plastic structures with non-rigid behaviours -- as do many large-scale natural phenomena, such as snowdrifts, oceans, and weather systems.

¹⁴ <https://en.wikipedia.org/wiki/Meccano>, <https://en.wikipedia.org/wiki/Tinkertoy> and <https://en.wikipedia.org/wiki/Lego>

Schrödinger's 1944 book stressed that quantum mechanisms can explain the structural stability of individual molecules. Quantum mechanisms can also explain how a set of atoms in different arrangements can form discrete stable structures with very different properties (e.g. in propane and isopropane, only the location of the single oxygen atom differs, but that alters both the topology and the chemical properties of the molecule)¹⁵. He also pointed out the relationship between the number of discrete changeable elements and information capacity, anticipating [Shannon 1948].

¹⁵ E.g. see James Ashenhurst's tutorial:
<http://www.masterorganicchemistry.com/2011/11/10/dont-be-futyl-learn-the-butyls/>

(While Shannon and most other thinkers who write about information treat information as essentially something to be *communicated*, Schrödinger (I think) understood that *use* of information is the key to its importance. Why communicate something, unless it is something that can be used? This shift of emphasis is crucial for understanding the many roles of information in biology. For example, perceptual information may or may not be communicated, but its primary role in an organism is its *use*, e.g. in controlling actions, or suggesting possible actions (among many other uses).

Some complex molecules with quantum-based structural stability are simultaneously capable of *continuous* deformations, e.g. folding, twisting, coming together, moving apart, etc., all essential for the role of DNA and other molecules in reproduction, and many other biochemical processes. This combination of discrete topological structure (forms of connectivity), used for storing very precise information for extended periods, and non-discrete spatial flexibility, used in assembling, replicating and extracting information from large structures, is unlike anything found in digital computers, although it can to some extent be approximated in digital computer models of molecular processes.

Highly deterministic, very small-scale, discrete interactions between very complex, multi-stable, enduring molecular structures, combined with continuous deformations (folding, etc.) that alter opportunities for the discrete interactions, may have hitherto unnoticed roles in brain functions, in addition to their profound importance for reproduction and growth. Much recent AI and neuroscience uses statistical properties of complex systems with many continuous scalar quantities changing randomly in parallel, unlike the discrete symbolic mechanisms used in logical and symbolic AI. Both appear to be far too restricted to model animal minds, but only when a wide variety of animal competences are taken into account. Selecting only simpler challenges can give the illusion of steady progress towards the ultimate goal of full replication of natural intelligence.

The Meta-Morphogenesis project extends the set of examples studied four decades earlier (e.g. in [Sloman 1978]) of types of mathematical discovery and reasoning that use perceived *possibilities* and *impossibilities* for change in geometrical and topological structures. Further work along these lines may help to reveal biological mechanisms that enabled the great discoveries by Euclid and

his predecessors that are still unmatched by AI theorem provers (discussed in Section 5).

IJCAI 2017 presentation (Added 29 Aug 2017)

An invited talk presented remotely at the IJCAI 2017 conference in Melbourne develops some of these ideas in a video presentation and accompanying web page, which overlaps partly with this one.

Why can't (current) machines reason like Euclid or even human toddlers?
(And many other intelligent animals)

Video <http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17>

Web page <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html>

2.6 Thermodynamic issues

The question sometimes arises whether formation of life from non-living matter violates the second law of thermodynamics, because life increases the amount of order or structure in the physical matter on the planet, reducing entropy. The standard answer is that the law is applicable only to closed systems, and the earth is not a closed system, since it is constantly affected by solar and other forms of radiation, asteroid impacts, and other external influences. The law implies only that our planet could not have generated life forms without energy from non-living sources, e.g. the sun (though future technologies may reduce or remove such dependence). Some of the ways in which pre-existing dispositions can harness external sources of energy to increase local structure are discussed in a collection of thoughts on entropy, evolution, and construction kits:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html>^[16]

¹⁶ Partly inspired by memories of a talk by Lionel Penrose in Oxford around 1960 about devices he called *droguli* - singular *drogulus*.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/entropy-evolution.html#new-droguli>

Such naturally occurring multi-stable physical structures seem to me to render redundant the apparatus proposed in [Deacon 2011] to explain how life apparently goes against the second law of thermodynamics. See https://en.wikipedia.org/wiki/Incomplete_Nature

Our discussion so far suggests that the FCK has two sorts of component: (a) a generic framework including space-time and generic constraints on what can happen in that framework, and (b) components that can be non-uniformly and dynamically distributed in the framework. The combination makes possible formation of galaxies, stars, clouds of dust, planets, asteroids, and many other lifeless entities, as well as supporting forms of life based on derived construction kits (DCKs) that exist only in special conditions. Some local conditions e.g. extremely high pressures, temperatures, and gravitational fields, (among others) can mask some parts of the FCK, i.e. prevent them from functioning. So, even if all sub-atomic particles required for earthly life exist at the centre of the sun, local factors can rule out earth-like life forms. Moreover, if the earth had been formed from a cloud of particles containing no carbon, oxygen, nitrogen, iron, etc., then no DCK able to support life as we know it could have emerged, since that requires a region of space-time with a specific manifestation of the FCK, embedded in a larger region that can contribute additional energy (e.g. solar radiation) and possibly other resources.

As the earth formed, new physical conditions created new DCKs that made the earliest life forms possible. [Ganti 2003], usefully summarised in [Korthof 2003] and [Fernando 2008], presents an analysis of requirements for a minimal life form, the "chemoton", with self-maintenance and reproductive capabilities. Perhaps still unknown DCKs made possible formation of pre-biotic

chemical structures, and also the *environments* in which a chemoton-like entity could survive and reproduce. Later, conditions changed in ways that supported more complex life forms, e.g. oxygen-breathing forms. Perhaps attempts to identify the first life form in order to show how it could be produced by the FCK are misguided, because *several* important pre-life construction kits were necessary: i.e. several DCKs made possible by conditions on earth were necessary for precursors. Some of the components of the DCKs may have been more complex than their living products, including components providing *scaffolding* for constructing life forms, rather than materials.

2.7 Scaffolding in construction kits

An important feature of some construction kits is that they contain parts that are used during assembly of products of the kit, but are not included in the products. For example, Meccano kits come with spanners and screwdrivers, used for manipulating screws and nuts during assembly and disassembly, though they are not normally included in the models constructed. Similarly, kits for making paper dolls and their clothing¹⁷ may include pencils and scissors, used for preparing patterns and cutting them out. But the pencils and scissors are not parts of the dolls or their clothing. When houses are built many items are used that are not part of the completed house, including tools and scaffolding frameworks to support incomplete structures.

¹⁷ https://en.wikipedia.org/wiki/Paper_doll

A loose analogy can be made with the structures used by climbing plants, e.g. rock-faces, trees, or frames provided by humans: these are essential for the plants to grow to the heights they need but are not parts of the plant. More subtly, rooted plants that grow vertically make considerable use of the soil penetrated by their roots to provide not only nutrients but also the stability that makes tall stalks or trunks possible, including in some cases the ability to resist strong winds most of the time. The soil forms part of the scaffolding.

A mammal uses parts of its mother as temporary scaffolding while developing in the womb, and continues to use the mother during suckling and later when fed portions of prey caught by parents. Other species use eggs with protective shells and food stores. Plants that depend on insects for fertilization can be thought of as using scaffolding in a more general sense.

This concept of scaffolding may be crucial for research into origins of life. As far as I know, nobody has found candidate non-living chemical substances made available by the FCK that have the ability spontaneously to assemble themselves into primitive life forms. It is possible that the search is doomed to fail because there never were such substances. Perhaps the earliest life forms required not only materials but also scaffolding - e.g. in the form of complex molecules that did not form parts of the earliest organisms but played an essential causal role in assembly processes, bringing together the chemicals needed by the simplest organisms. Perhaps at some stage enormous gravitational fields were required that could not be reproduced in our laboratories.

Evolution might later have produced new organisms without that reliance on the original scaffolding. The scaffolding mechanisms might later have ceased to exist on earth, e.g. because they were consumed and wiped out by the new life forms, or because physical conditions changed that prevented them forming but did not destroy the newly independent organisms. Or the scaffolding might have existed only in another part of the universe, whose products were later dispersed. Similar suggestions are made in [[Mathis, Bhattacharya, Walker. 2015](#)] - and for all I know they have also been made elsewhere.

It is quite possible that many evolutionary transitions, including transitions in information processing, our main concern, depended on forms of scaffolding that later did not survive and were no longer needed to maintain what they had helped to produce. So research into evolution of information processing, our main goal, is inherently partly speculative.

2.8 Biological construction kits

How did the FCK generate complex life forms? Is the Darwin-Wallace theory of natural selection the whole answer, as suggested in [Bell 2008]? "*Living complexity cannot be explained except through selection and does not require any other category of explanation whatsoever.*" No: the explanation must include both *selection* mechanisms and *generative* mechanisms, without which selection processes will not have a supply of new viable options. Moreover, insofar as environments providing opportunities, challenges and threats are part of the selection process, the construction kits used by evolution include mechanisms not intrinsically concerned with life, e.g. volcanoes, earthquakes, asteroid impacts, lunar and solar tides, and many more.

The idea of evolution producing construction kits is not new, though they are often referred to as "toolkits". [Coates, Umm-E-Aiman, Charrier. 2014] ask whether there is "a genetic toolkit for multicellularity" used by complex life-forms. (I use the label "construction kit" to refer to the combination of tools and the materials on which the tools operate.) 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 generative mechanisms and selection mechanisms change during evolution. 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 have also occurred in human-designed computing mechanisms, including new forms of hardware, new forms of virtual machinery, and networked social systems all unimagined by early hardware designers. Similar changes during evolution produced new biological construction kits, e.g. grammars, planners, geometrical constructors, not well understood by thinkers familiar only with physics, chemistry and numerical mathematics.

Biological DCKs produce 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 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 also used unwittingly in practical activities.¹⁸

¹⁸ Such as putting a shirt on a child (I think Piaget noticed some of the requirements): <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/shirt.html>.

I suspect many animals come close to this in their *systematic* but unconscious abilities to perform complex actions that use mathematical features of environments. Abilities used unconsciously in building nests or in hunting and consuming prey may overlap with topological and geometrical competences of human mathematicians. (See Section 6.2). E.g. search for videos of weaver birds building nests.

2.9 Cognition echoes reality (INCOMPLETE DRAFT. Added 13 Apr 2016)

Each new construction kit produced either by evolution or by environmental changes (e.g. geological processes) or by human technology, makes possible new products that in the presence of life can provide opportunities for evolution to produce or extend construction kits for building information processing systems (e.g. cognitive mechanisms). Then users of the new information processing abilities may be better able than their predecessors to cope with the new opportunities, impediments, constraints, dangers, etc.

These are deeper changes than merely coping with new products of old construction kits, since those products will involve old structures, properties and relationships for which previously evolved information processing capabilities may suffice. Thinking about how to use newly manufactured slightly bigger bricks may not require major new design capabilities. In contrast, when steel becomes available as a building material for the first time, bigger changes in architectural design capabilities may be required before the new potential can be realised. This in turn may require a new construction kit for building information processing systems extending minds of designers. In this case the new construction kit may be an extended form of education in architectural design.

Later, some of the products of the new construction kit may be combined either with old pre-existing structures, or with independently evolved new structures, producing yet another new construction kit providing opportunities for yet another layer of complexity in information processing mechanisms -- an opportunity for a new or modified construction kit for building information processing mechanisms.

So there is a potentially endless collection of new (branching) layers of complexity both in information processing systems and in entities about which information can be acquired and used. (This section may be expanded, modified, or moved, later.)

2.10 Goodwin's "Laws of Form": Evolution as a form-maker and form-user DRAFT Added: 18 Apr 2016

An old idea in biology is that besides the operation of inherited biological mechanisms the products of evolution also result in part from "Laws of form" that constrain what can evolve and develop.

On Wed 19th March 2014 a wonderful talk was given to the Birmingham Mathematics colloquium by [Professor Reidun Twarock](#) (York University) on "Viruses and geometry - a new perspective on virus assembly and anti-viral therapy". The abstract stated:

"A large number of human, animal and plant viruses have protein containers that provide protection for their genomes. In many cases, these containers, called capsids, exhibit symmetry, and they can therefore be modelled using techniques from group, graph and tiling theory. It has previously been assumed that their formation from the constituent protein building blocks can be fully understood as a self-assembly process in which viral genomes are only passive passengers. Our mathematical approach, in concert with techniques from bioinformatics, biophysics and experiment, provides a new perspective: It shows that, by contrast, interactions between viral genome and capsid play vital cooperative roles in this

process in the case of RNA viruses, enhancing assembly efficiency and fidelity. We use the graph theoretical concept of Hamiltonian path to quantify the resulting complexity reduction in the number of assembly pathways, and discuss implications of these insights for a novel form of anti-viral therapy."

An implication seems to be that the reproduction process for such a virus does not require all details to be specified by the genome because the symmetries in the construction constrain possible ways in which the components can be assembled, enormously reducing the number of decisions that need to be made about where molecules should go. But those symmetries (control by form) are also not enough: it's a collaborative process. In the context of this paper the "laws of form" need to be understood as partly referring to potential and limitations of currently available construction kits. So evolution can change the laws of form which in turn help to change later forms of evolution.

This seems to be related to some of the ideas of [Brian Goodwin](#) about laws of form [[Goodwin 1995](#)] that enable and constrain products of biological evolution, rather than everything being a result of fitness requirements shaping an amorphous space of possibilities. Some of those ideas come from earlier work by [D'Arcy Thompson](#), Goethe and others. See [Boden \(2006\) Sections 15x\(b-d\), Vol 2](#), for more on this.

I want to go further and suggest that the processes of natural selection often discover mathematical domains that can be put to use not only during morphogenesis/epigenesis, but also during the ongoing functioning of individual organisms, under the (partial) influence of information-based control mechanisms. A simple example that is re-used in many contexts is the mathematics of homeostasis (control by negative feedback loops). There are probably many more examples waiting to be discovered, including the mechanisms that first enabled humans to discover and prove theorems, e.g. in geometry and arithmetic (long before geometry had been mapped into arithmetic and algebra, and long before the development of the axiomatic method based on logic).

So, whereas Goodwin and others have thought of "Laws of form" as controlling what Goodwin in 1995 described as "*a dance through morphospace, the space of the forms of organisms*", we need to expand his idea to laws of form that enable and control the space of types of information processing system and their behaviours. These laws can't be modelled using the mathematics of dynamical systems with a large, but fixed, collection of numerical parameters, defining possible trajectories through state space. Turing's paper on chemistry-based morphogenesis [[Turing 1952](#)] was an early pioneering demonstration of some types of "dance through morphospace". But if Turing had not died only two years after it was published, it seems likely that he would have combined those ideas with his earlier, better known ideas about forms of information processing. Perhaps he would have produced a much deeper richer theory of construction kits, many of them products of the dances of older kits, than I can.

See also <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/mathsem.html> How could evolution produce mathematicians from a cloud of cosmic dust?

[CONTENTS](#)

3 Concrete (physical), abstract and hybrid construction kits

Products of a construction kit may be concrete, i.e. physical, or abstract, like a theorem, a sentence, or a symphony; or hybrid, e.g. a written presentation of a theorem 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 also have new shapes and new behaviours. Those are *concrete* construction kits. The FCK is (arguably?) a concrete construction kit. Lego, Meccano, twigs, mud, and stones, can all be used in construction kits whose constructs are physical objects occupying space and time: *concrete* construction kits.

Abstract kits: There are also non-spatial *abstract* construction kits, 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. Using an abstract construction kit, e.g. doing mental arithmetic or composing poetry in your head, requires use of one or more physical construction kits, that directly or indirectly implement features of the abstract kit.

There are (deeply confused) fashions emphasising "embodied cognition" and "symbol grounding" (previously known as "concept empiricism" and demolished by Immanuel Kant and 20th Century philosophers of science). These fashions disregard many examples of thinking, perceiving, reasoning and planning, that require abstract construction kits. For example, planning a journey to a conference does not require physically trying possible actions, like water finding a route to the sea. Instead, one can use an abstract construction kit able to *represent* possible options and ways of combining them.

Being able to talk requires use of a grammar specifying 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 information. Sentences allowed by a grammar for English are abstract objects that can be instantiated physically in written text, printed text, spoken sounds, morse code, etc.: so a grammar is an abstract construction kit whose constructs can have concrete (physical) instances. Moreover, it is clear from many computer-based systems, including email systems, search engines, online reference materials, chat systems, that it is also possible for words, phrases, sentences, paragraphs, etc. in human languages to be instantiated in *virtual* machinery running on computers, or networks of computers.

The idea of a grammar is not restricted to verbal forms: it can be extended to many complex structures, e.g. grammars for sign languages, circuit diagrams, maps, proofs, architectural layouts and even molecules, as noticed by AI vision researchers in the 1960s. Some relevant publications were assembled in [\[Kaneff 1970\]](#)

A grammar does not fully specify a language: a structurally related *semantic* construction kit is required for building possible *meanings*. Use of a language depends on language users (such as information processing subsystems), for which more complex construction kits are required, including products of evolution, development and learning. Evolution of various types of language is discussed in [\[Sloman 2008/2011\]](#). It is argued there that not only humans, but also many other intelligent species, and pre-verbal human infants and toddlers, require languages used for *internal* functions (e.g. perception, intention formation, planning, reasoning and learning), not *external*

communication between individual organisms.

In computers, physical mechanisms implement abstract construction kits via intermediate abstract kits - virtual machines; and presumably something similar also happens in brains, using types of virtual machinery discovered long ago by evolution.

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 objects (e.g. blocks) to instantiate arithmetic, which may lead to the discovery of prime numbers when certain attempts at rearrangement fail - and an explanation of the impossibility is found.¹⁹

¹⁹ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#primes>

In some hybrid construction kits, such as games like chess, the concrete (physical) component may be redundant for some players: e.g. chess experts who can play without physical pieces on a board. But communication of moves needs physical mechanisms, as does the expert's brain (in ways that are not yet understood). Related abstract structures, states and processes can also be implemented in computers, which can now play chess better than most humans, without replicating human brain mechanisms. In contrast, physical components are indispensable in hybrid construction kits for outdoor games, like cricket [Wilson 2015]. (I don't expect to see good robot cricketers soon.)

Physical computers, programming languages, operating systems and virtual machines form hybrid construction kits that make things happen when they run.

Logical systems with axioms and inference rules can be thought of as abstract kits supporting construction of logical proof-sequences, usually combined with a physical notation for written proofs.

Purely logical systems are mathematical abstractions which, like the number 99, cannot have physical causal powers, whereas concrete instances can, e.g. teaching a student, or programming a computer, to distinguish valid and invalid proofs. Natural selection "discovered" the power of hybrid construction kits using virtual machinery long before human engineers did.

In particular, biological virtual machines used by animal minds outperform current engineering designs in some ways, but they also generate much confusion in the minds of philosophical individuals who are aware that something more than purely physical machinery is at work, but don't yet understand how to implement virtual machines in physical machines [Sloman Chrisley 2003, Sloman 2010, Sloman 2013a].

Animal perception, learning, reasoning, and intelligent behaviour require *hybrid* construction kits. Scientific study of such kits is still in its infancy. Work done so far on the Meta-Morphogenesis project suggests that natural selection "discovered" and used a staggering variety of types of hybrid construction kit (combining collections of virtual machines interacting with one another and with physical machines, some parts of which are controlled by virtual machines).

Thinking processes that produce physical sounds (spoken sentences) share this feature with virtual machines in computers, e.g. email systems, that produce textual displays on a screen.

There is still much to be learned about hybrid biological construction kits that are essential for reproduction, for developmental processes (including physical development and learning), for performing complex behaviours, for making mathematical, scientific and technological discoveries, and for social/cultural phenomena.

Some, like flocking/murmurations of starlings, are relatively simple to explain (at a high level of abstraction, ignoring most of the biological detail):

<https://www.youtube.com/watch?v=8vhE8ScWe7w>

Modelling in detail the social life of a species of apes is a far more ambitious goal, that may be unattainable for a very long time, especially if the apes are humans.

3.1 Kits providing external sensors and motors

Some construction kits can be used to make toys with moving parts, e.g. wheels or grippers, that interact with the environment. A toy car may include a spring, whose potential energy can be transformed into mechanical energy via gears, axles and wheels in contact with external surfaces. Further interactions, altering the direction of motion, may result from collisions with fixed or mobile objects in the environment or may be produced by some control device. More sophisticated kits have sensors that provide information for internal mechanisms that use information to select options for movement, etc. More sophisticated kits combine several sensors, e.g. optical, auditory, tactile, haptic, proprioceptive, inertial, or chemical sensors, that provide information that internal mechanisms can use to evaluate and select goals or options for achieving existing goals, control actions, interact with conspecifics, predict events in the environment, evaluate hypotheses, and other functions. For example, in rats, whiskers gain information allowing frequent changes of direction and speed of motion, among other things. E.g. see <http://www.ratbehavior.org/RatWhiskers.htm>

As noted in [Sloman 1978,Ch.6], the distinction between internal and external components is often arbitrary - a fact frequently rediscovered. For example, a music box may perform a tune under the control of a rotating disc with holes or spikes. The disc can be thought of as part of the music box, or as part of a changing environment.

If a toy train set has rails or tracks used to guide the motion of the train, then the wheels can be thought of as sensing the environment and causing changes of direction. 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 information and the forces required to change direction. A 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 wolf following a scent trail left by its prey. An unseen wall uses force to stop your forward motion in a dark room, whereas a visible, or lightly touched, wall provides only information [Sloman 2011].

3.2 Mechanisms for storing, transforming and using information

Often information is acquired, used, then lost because it is 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, change direction, etc. In more complex control systems, with "offline intelligence" sensor information is saved, possibly combined with previously stored information, and remains available for use on different occasions for different purposes.²⁰

²⁰ [Trehub 1991] proposed an architecture for vision that allows snapshots from visual saccades to be integrated in a multi-layer fixation-independent visual memory.

In the "offline" 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. A control decision at one time may need 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. Just as engineers sometimes need to recalibrate equipment evolution can provide mechanisms for testing and adjusting such sensors, e.g. by altering pupil size when the level of ambient illumination changes.

Similarly, when used offline, some time after acquisition, stored information may need to be checked for minor inaccuracy, or complete falsity caused by the environment changing, as opposed to sensor faults. The difference between offline and online uses of visual information has caused much confusion among researchers, including muddled attempts to interpret the difference in terms of "what" and "where" information.²¹ Contrast [Sloman 1983].

²¹ http://en.wikipedia.org/wiki/Two-streams_hypothesis

As organisms moved around in larger portions of the environment with different resources, opportunities, and dangers at different locations, this produced increasing requirements for items of stored information to be used in different ways, for different purposes, at different times and locations from the time and location of acquisition. In short: evolution of more complex competences produced increasingly disembodied forms of information processing, a fact that seems not to be understood by rigid defenders of "embodied cognition", sometimes referred to as "enactive cognition". Fortunately, many thinkers have seen through this confusion, including Rescorla [2015].

Many ways of acquiring and using information have been discovered and modelled by AI researchers, psychologists, neuroscientists, biologists and others. However, evolution produced many more. Some of them require not just additional storage space but also very different sorts of information processing architecture, whereas AI engineers typically seek one architecture for a project.

Varieties of biologically useful, biologically inspired, architecture are explored in papers and presentations in the Birmingham Cognition and Affect project [CogAff (1991-...)], including [Sloman 1983], [Sloman 1993], [Sloman 2003], [Sloman 2006], [Sloman 2013a]. However, I make no claim regarding completeness: I suspect there are many computational designs produced by evolution that we have not yet thought of, some of which may still be important in human brains, e.g. below the level of neurons.

A complex biological architecture may use sub-architectures that evolved at different times, meeting different needs in different niches. In particular, I suspect there are biological mechanisms for handling vast amounts of rapidly changing incoming information (visual, auditory, tactile, haptic, proprioceptive, vestibular) by using several different sorts of short-term storage plus processing subsystems, operating on different time-scales in parallel, including a-modal information structures. This undermines any theory that assumes there is one short-term memory and one long-term memory.

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 species that produce new subsystems, or modify old ones, during individual development, e.g. during different phases of learning by apes, elephants, and humans, as described in Section 2.3, contradicting the common assumption that a computational architecture must remain fixed.²²

²² The BICA society aims to bring together researchers on biologically inspired cognitive architectures. Some examples are here: <http://bicasociety.org/cogarch/>

3.3 Mechanisms for controlling position, motion and timing

All concrete construction kits (and some hybrid kits) share a deep common feature insofar as their components, their constructs and their construction processes involve space and time, both during assembly and while working. Those behaviours include both relative motion of parts, 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 spatiality is that objects built from different construction kits can interact, by changing their spatial relationships (e.g. if one object, or part of one object, enters, encircles or grasps another), applying forces transmitted through space, and using spatial sensors to gain information used in control.

Products of different kits can interact in varied ways, e.g. one being used to assemble or manipulate another, or one providing energy or information for the other. Contrast the problems of getting software components available on a computer to interact sensibly: merely locating them in the same virtual or physical machine will not suffice. Some rule-based systems are composed of condition-action rules, managed by an interpreter that constantly checks for satisfaction of conditions. Newly added rules may then be invoked simply because their conditions become satisfied, though "conflict resolution" mechanisms may be required if the conditions of more than one rule are satisfied.²³

²³ Our SimAgent toolkit is an example [Sloman 1996b].
<http://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html>

New concrete kits can be formed by combining two or more kits. In some cases this will require modification of a kit, e.g. combining Lego and Meccano by adding pieces with Lego studs or holes alongside Meccano sized screw holes. In other cases mere spatial proximity and contact suffices, e.g. when one construction kit is used to build a platform and others to assemble a house on it. Products of different biological construction kits may also use complex mixtures of juxtaposition and adaptation.

Objects that exist in space and time often need timing mechanisms. Organisms use "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 and adjustable rates of change: e.g. a bird in flight approaching a perch; an animal running to escape a predator and having to decelerate as it approaches a tree it needs to climb; 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 suitable locations relative to the intended grasping points on the object. (This can be very difficult for robots, when grasping novel objects in novel situations: if they use ontologies that are too simple.) There are also biological mechanisms for controlling or varying rates of production of chemicals (e.g. hormones).

So biological construction kits need many mechanisms able to measure time intervals and to control rates of repetition or rates of change of parts of the organism. These kits may be combined with other sorts of construction kit that combine temporal and spatial control, e.g. changing speed and direction simultaneously, or altering grasp width while reducing distance from grasper to thing grasped (moving fingers and thumb further apart as your hand approaches a mug to be grasped).

A violinist has to learn to control several kinds of change in relationships of a violin bow simultaneously, including continuous and discrete changes. The continuous changes include changes of speed of motion of the bow-hairs across one or more violin strings, changes of pressure of the bow on the string, changes of location on a string where the bow is in contact with the string (i.e. nearer or further from the violin bridge). In other contexts there are changes in frequency of direction change (moving the bow up and down more or less quickly) changes in whether the bow moves constantly in contact with the string (legato) or while bouncing on the string (staccato). In addition there are parallel changes of various sorts in the other hand (usually left hand) as the fingers move against the strings, changing location of contact, changing pressure, changing between continuous (glissando) and discontinuous contact, and whether only one string is being used at a time or more than one.

The violin examples are highly specialised and usually require explicit teaching by an expert and years of practice. However everyday life contains many other examples of simultaneous change, including the changes of relationships between tongue, lips, cheeks, and teeth in the mouth cavity during consumption of different kinds of food, or while speaking. The latter also requires complex coordination of breathing and control of vocal chords.

For all of these various kinds of control functions specific biological mechanisms are required, including control mechanisms as well as controlled mechanisms. And the construction of those mechanisms during development and growth of individual organisms requires the use of construction kits of various kinds whose existence is not as obvious as what they construct. I suspect that our current knowledge of the variety of construction kits, how they work, how they develop in individuals, how they relate to genetic mechanisms, and how they evolved, is still miniscule, and one of the consequences may be that many of the functions for which such mechanisms have been used have not been noticed or have not been studied in depth.

In consequence, our understanding of the variety of requirements, especially information processing requirements, for mechanisms performing similar functions in intelligent machines, including autonomous robots, may be very limited, contributing to serious limitations of current

intelligent machines. Those deficiencies in information processing abilities and control functions will not be removed simply by increasing the physical similarities between robot bodies and animal bodies, as sometimes implied by "embodiment" enthusiasts.

The deficiencies will be partly due to our lack of understanding of requirements for construction kits used in creating the mechanisms during development and learning.

3.4 Combining construction kits

At the molecular level there is now a vast, and rapidly growing, amount of biological 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.

In computers, 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 cases most people encounter when learning arithmetic. A function could be a compiler, its arguments could be arbitrarily complex programs in a high level programming language, and the outputs of the function might be either a report on syntactic errors in the input program, or a machine code program ready to run.

Applying functions to arguments is very different from assembling structures in space-time, where inputs to the process form parts of the output. If computers are connected via digital to analog interfaces, linking them to 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.

The additional richness is present in the modes of interaction of chemical structures that include both digital (on/off chemical bonds) and continuous changes in relationships, as discussed in [\[Turing 1952\]](#), the paper on chemistry-based morphogenesis that inspired this Meta-Morphogenesis project [\[Sloman 2013b\]](#).

3.5 Combining abstract construction kits

Section [2.1](#) showed how a new DCK using combinations of old components can make some new developments very much quicker to reach -- fewer steps are required, and the search space for a sequence of steps to a solution may be dramatically reduced.

Combining *concrete* construction kits uses space-time occupancy. Combining *abstract* construction kits is less straightforward. Sets of letters and numerals are combined to form labels for chess board squares, e.g. "a2", "c5", etc. A human language and a musical notation can form a hybrid system for writing songs. A computer operating system (e.g. Linux) can be combined with programming languages (e.g. Lisp, Java).

In organisms, as in computers, products of different kits may share *information*, e.g. information for sensing, predicting, explaining or controlling, including information about information [\[Sloman 2011\]](#).

Engineers combining different kinds of functionality often find it useful to design re-usable information processing *architectures* that provide frameworks for combining different mechanisms and information stores^[22], especially in large projects where different teams work on sensors, learning, motor systems, reasoning systems, motivational systems, various kinds of metacognition, etc., using different specialised tools. The toolkit mentioned in Note [23](#) is an example framework.

It is often necessary to support different sorts of **virtual** machinery interacting simultaneously with one another and with internal and external physical environments, during perception and motion. This may require new general frameworks for assembling complex *information processing architectures*, accommodating multiple interacting virtual machines, with different modifications developed at different times [[Minsky 1987](#),[Minsky 2006](#)] [[Sloman 2003](#)].

Self-extension is a topic for further research.^[24], [[Sloman 2009](#)].

²⁴ As discussed in connection with "toddler theorems" in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>
Contributions from observant parents and child-minders are welcome. I think deeper insights come from extended individual developmental trajectories than from statistical snapshots of many individuals.

Creation of new construction kits may start by simply recording parts of successful assemblies, or, better still, parametrised parts, so that they can easily be reproduced in modified forms - e.g. as required for organisms that change size and shape while developing.

Eventually, parametrised stored designs may be combined to form a "*meta-construction kit*" able to extend, modify or combine previously created construction kits -- as human engineers have recently learnt to do in debugging toolkits. Evolution needs to be able to create new meta-construction kits using natural selection. Natural selection, the great creator/meta-creator, is now spectacularly aided and abetted by its products, especially humans!

[CONTENTS](#)

4 Construction kits generate possibilities and impossibilities

Explanations of how things are possible (Sect. [1](#)) can refer to construction kits, either manufactured, e.g. Meccano and Lego, or composed of naturally occurring components, e.g. boulders, mud, or sand. (Not all construction kits have sharp boundaries.) Each kit makes possible certain types of construct, instances of which can be built by assembling parts from the kit. Some construction kits use *products of products of biological evolution*, e.g. birds' nests assembled from twigs or leaves.

In some kits, features of components, such as shape, are inherited by constructed objects. E.g. objects composed only of Lego bricks joined in the "standard" way have external surfaces that are divisible into faces parallel to the surfaces of the first brick used. However, if two Lego bricks are joined at a corner only, using only one stud and one socket, it is possible to have continuous relative rotation (because studs and sockets are circular), violating that constraint, as Ron Chrisley pointed out in a conversation.

This illustrates the fact that constructed objects can have "emergent" features none of the components have, e.g. a hinge is a non-rigid object that can be made from two rigid objects with aligned holes through which a screw is passed.

So, 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.

4.1 Construction kits for making information-users

Not everything that can play a role in acquisition, storage or transfer of information has information processing capabilities. Consider a lump of plasticine or damp clay that can be deformed under pressure, then retains the deformation. If a coin is pressed against it the lump will change its shape.

Entities with information processing capabilities (e.g. archaeologists) can use the depression as a source of information about the coin. But the deformed lump of material is not an information user. If the depression is used to control a process, e.g. making copies of the coin, or to help a historian years later, then the deformed material is used as a source of information about the coin.

The fact that some part of a brain is changed by perceptual processes in an organism does not imply that that portion of the brain is an information user. It may play a role analogous to the lump of clay, or a footprint in soil. Additional mechanisms are required if the information is to be *used*: different mechanisms for different types of use.

A photocopier acquires information from a sheet of paper, but all it can do with the information is produce a replica (possibly after slight modifications such as changes in contrast, intensity or magnification). Different mechanisms are required for recognising text, correcting spelling, analysing the structure of an image, interpreting it as a picture of a 3-D scene, using information about the scene to guide a robot, building a copy of the scene, or answering a question about which changes are possible.

Thinking up ways of using the impression of a coin as a source of information about the coin is left as an exercise for the reader.

Biological construction kits for producing information processing mechanisms evolved at different times. [Sloman 1993] discusses the diversity of uses of information from biological sensors, including sharing of sensor information between different uses, either concurrently or sequentially. Some of the mechanisms use intermediaries, such as sound or light, to gain information about the source or reflector of the sound or light. This can be used in taking decisions, e.g. whether to flee, or used in controlling actions, such as grasping or walking towards a source of information.

Some mechanisms that use information seem to be direct products of biological evolution, such as mechanisms that control reflex protective blinking. Others are grown in individuals by epigenetic mechanisms influenced by context, as explained in Section 2.3. For example, humans in different cultures start with a generic language construction kit (sometimes misleadingly labelled a "universal grammar") which is extended and modified to produce locally useful language understanding/generating mechanisms.

Language-specific mechanisms --- such as mechanisms for acquiring, producing, understanding and correcting textual information --- evolved long after mechanisms able to use visual information for avoiding obstacles or grasping objects, shared between far more types of animal.

In some species there may be diversity in the construction kits produced by individual genomes, leading to even greater diversity in adults, if they develop in different physical and cultural environments using epigenetic mechanisms discussed above.

4.2 Different roles for information

Despite huge diversity in biological construction kits and the mechanisms in individual organisms, some themes recur, such as functions of different sorts of information in control: e.g. information about how things actually are or might be ("belief-like" information contents), information about how things need to be or might need to be for the individual information user ("desire-like" information contents), and information about how to achieve or avoid certain states ("procedural" information contents).

Each type of role has different subtypes, differing across species, across members of a species and across developmental stages in an individual. How a biological construction kit supports the requirements for all the subtypes depends on the environment, the animal's sensors, its needs, the local opportunities, and the individual's history.

Different mechanisms performing such a function may share a common evolutionary precursor after which they diverged. Moreover, mechanisms with similar functions can evolve independently - convergent evolution.

Information relating to targets and how to achieve or maintain them is *control* information: the most basic type of biological information, from which all others are derived. A simple case is a thermostatic control, discussed in [McCarthy 1979]. It has two sorts of information: (a) a target temperature (desire-like information) (b) current temperature (belief-like information). A discrepancy between them causes the thermostat to select between turning a heater on, or off, or doing nothing.

This very simple homeostatic mechanism uses **information** and an **energy source** to achieve or maintain a state. There are very many variants on this schema, according to the type of target (e.g. a measured state or some complex relationship) the type of control (on, off, or variable, with single or multiple effectors), and the mechanisms by which control actions are selected, which may be modified by learning, and may use simple actions or complex plans.

As [Gibson 1966] pointed out, acquisition of information often requires cooperation between processes of sensing and acting. Saccades are visual actions that constantly select new information samples from the environment (or the optic cone). Uses of the information vary widely according to context, e.g. controlling grasping, controlling preparation for a jump, controlling avoidance actions, or sampling text to be read. A particular sensor can therefore be shared between many control subsystems [Sloman 1993], and the significance of the sensor state will depend partly on which subsystems are connected to the sensor at the time, and partly on which other mechanisms receive information from the sensor (which may change dynamically - a possible cause of some types of "change blindness").

The study of varieties of use of information in organisms is exploding, and includes many mechanisms on molecular scales as well as many intermediate levels of informed control, including sub-cellular levels (e.g. metabolism), physiological processes of breathing, temperature maintenance, digestion, blood circulation, control of locomotion, feeding and mating of large animals and coordination in communities, such as collaborative foraging in insects and trading systems of humans.

In particular, slime moulds include spectacular examples in which modes of acquisition and use of information by an individual change dramatically.²⁵

²⁵ <http://www.theguardian.com/cities/2014/feb/18/slime-mould-rail-road-transport-routes>




Figure 4: Evolutionary transitions from molecules to intelligent animals

Between the simplest and most sophisticated organisms there are many intermediate forms with very different information processing requirements and capabilities.

The earliest organisms must have acquired and used information about things inside themselves and in their immediate vicinity, e.g. using chemical detectors in an enclosing membrane.

Later, evolution extended those capabilities in dramatic ways (crudely indicated in Fig. 4). In the simplest cases, local information is used immediately to select between alternative possible actions, as in a heating control, or a trail-following mechanism.

Uses of motion in haptic and tactile sensing and use of saccades, changing vergence, and other movements in visual perception, all use the interplay between sensing and doing, in "online intelligence".

But there are cases ignored by Gibson and anti-cognitivists, namely organisms that exhibit "offline intelligence", using perceptual information for tasks other than controlling immediate reactions. Examples include: reasoning about remote future possibilities, attempting to explain something observed, and working out that bending a straight piece of wire will enable a basket of food to be lifted out of a tube as illustrated in Fig. 4 [Weir, Chappell, Kacelnik . 2002].

Such uses of new information require the new information to be combined with previously acquired information about the environment, including particular information about individual objects and their locations or states, general information about learnt laws or correlations and information about what is and is not possible (Note [13]).

An information-bearing structure (e.g. the impression of a foot or the shape of a rock) can provide very different information to different information-users, or to the same individual at different times, depending on (a) what kinds of sensors they have, (b) what sorts of information processing (storing, analysing, comparing, combining, synthesizing, retrieving, deriving, using...) mechanisms they have, (c) what sorts of needs or goals they can serve by using various sorts of information (knowingly or not), and (d) what information they already have.

So, from the fact that changes in some portion of a brain correlate 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 - since typically that will depend on context: the same sensory input can be used in very different ways in different contexts. More details are provided in [\[Sloman 1993\]](#).

4.3 Motivational mechanisms

It is often assumed that every information user, *U*, constantly tries to achieve rewards or avoid punishments (negative rewards), and that each new item of information, *I*, will make some actions more likely for *U*, and others less likely, on the basis of what *U* has previously learnt about which actions increase positive rewards or decrease negative rewards under conditions indicated by *I*.

But animals are not restricted to acting on motives selected *by them* on the basis of expected rewards. They may also have motive generators that are simply triggered as "internal reflexes" just as evolution produces phototropic reactions in plants without giving plants any ability to anticipate benefits to be gained from light.

Some reflexes, instead of directly triggering *behaviour*, trigger *new motives*, which may or may not lead to behaviour, depending on the importance of other competing motives. For example, a kind person watching someone fall may acquire a motive to rush to help - not acted on if competing motives are too strong. It is widely believed that all motivation is reward-based.

But a new motive triggered by an internal reflex need not be associated with some reward. It may be "architecture-based motivation" rather than "reward-based motivation" [\[Sloman 2009\]](#). Triggering of architecture-based motives in playful intelligent young animals can produce kinds of delayed learning that the individuals could not possibly anticipate [\[Karmiloff-Smith 1992\]](#).

Unforeseeable biological benefits of automatically triggered motives include acquisition of new information by sampling properties of the environment. The new information may not be immediately usable, but in combination with information acquired later and genetic tendencies activated later, as indicated in Fig. [3](#), it may turn out to be important, during hunting, caring for young, or learning a language.

A toddler may have no conception of the later potential uses of information gained in play, though the ancestors of that individual may have benefited from the presence of the information gathering reflexes. In humans this seems to be crucial for mathematical development.

During evolution, and also during individual development, the sensor mechanisms, the types of information processing, and the uses to which various types of information are put, become more diverse and more complex, while the information processing architectures allow more of the processes to occur in parallel (e.g. competing, collaborating, invoking, extending, recording, controlling, redirecting, enriching, training, abstracting, refuting, or terminating).

Without understanding how the architecture grows, which information processing functions it supports, and how they diversify and interact, we are likely to reach wrong conclusions about biological functions of the parts: e.g. over-simplifying the functions of sensory subsystems, or over-simplifying the variety of concurrent control mechanisms producing behaviours.

Moreover, the architectural knowledge about how such a system works, like information about the architecture of a computer operating system, may not be expressible in sets of equations, or statistical learning mechanisms and relationships. (Ideas about architectures for human information processing can be found in [[Simon 1967](#),[Minsky 1987](#),[Minsky 2006](#), [Laird, Newell, Rosenbloom . 1987](#),[Sloman 2003](#),[Sun 2006](#)], among many others.)

Construction kits for building information processing architectures, with multiple sensors and motor subsystems, in complex and varied environments, differ widely in the designs they can produce. Understanding that variety is not helped by disputes about which architecture is best. A more complete discussion would need to survey the design options and relate them to actual choices made by evolution or by individuals interacting with their environments.

[CONTENTS](#)

5 Mathematics: Some constructions exclude or necessitate others

Physical construction kits (e.g. Lego, plasticine, or a combination of paper, scissors and paste) have parts and materials 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, twisting and elastic or inelastic resistance to deformation).

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

Construction kits can also explain necessity and impossibility. A construction kit with a large initial set of generative powers can be used to build a structure realising some of the kit's possibilities, in which some further possibilities are excluded, namely all extensions that do not include what has so far been constructed. If a Meccano construction has two parts in a substructure that fixes them a certain distance apart, then no extension can include a new part that is wider than that distance in all dimensions and is in the gap. Some extensions to the part-built structure that were previously possible become impossible unless something is undone. That example involves a limit produced by a gap size. There are many more examples of impossibilities that arise from features of the construction kit.




Figure 5: Mary Pardoe's proof of the triangle sum theorem.

The sequence 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. It nowhere mentions Euclid's parallel axiom, used by "standard" proofs. This unpublished proof was reported to me by Mary Pardoe, a former student who became a mathematics teacher, in the early 1970s.

Euclidean geometry includes a construction kit that enables construction of closed planar polygons (triangles, quadrilaterals, pentagons, etc.), with interior angles whose sizes can be summed. If the polygon has three sides, i.e. it is a triangle, then the interior angles must add up to exactly half a rotation. Why?

In this case, no physical properties of a structure (e.g. rigidity or impenetrability of materials) are involved, only spatial relationships. Figure 5 provides one way to answer the question, unlike the standard proofs, which use parallel lines. It presents a proof, found by Mary Pardoe, that internal angles of a planar triangle sum to a straight line, or 180 degrees.

Most humans are able 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 a common feature. This may have evolved from earlier abilities to reason about changing affordances in the environment [Gibson 1979]. Current AI perceptual and reasoning systems still lack most of these abilities, and neuroscience cannot yet explain what's going on (as opposed to where it's going on?). (See Note [13]).

These illustrate 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 consequence: they make some further structures *possible* (e.g. constructing a circle that passes through all the vertices of the triangle), and other structures *impossible* (e.g. relocating the corners of the triangle so that the angles add up to 370 degrees).

These possibilities and impossibilities are *necessary* consequences of previous selection steps. The examples illustrate how a construction kit with mathematical relationships can provide the basis for necessary truths and necessary falsehoods in some constructions (as argued in [Sloman 1962, Chap 7]).²⁶

²⁶ Such relationships between possibilities provide a deeper, more natural, basis for understanding modality (necessity, possibility, impossibility) than so called "possible world semantics". I doubt that most normal humans who can think about possibilities and impossibilities base that on thinking about truth in the whole world, past, present and future, and in the set of alternative worlds.

Being able to think about and reason about alterations in some limited portion of the environment is a very common requirement for intelligent action [Sloman 1996a]. It seems to be partly shared with other intelligent species, e.g. squirrels, nest-builders, elephants, apes, etc. Since our examples of making things possible or impossible, or changing ranges of possibilities, are examples of causation (mathematical 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 5.3.²⁷

²⁷ 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.

Neuroscientific theories about information processing in brains currently seem to me to omit the processes involved in such mathematical discoveries, so AI researchers influenced too much by neuroscience may fail to replicate important brain functions. Progress may require major conceptual advances regarding what the problems are and what sorts of answer are relevant.

We now consider ways in which evolution itself can be understood as discovering mathematical proofs - proofs of possibilities.

5.1 Proof-like features of evolution

A subset of the FCK 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, and mathematical proofs. However, starting from those parts will make some designs unreachable except by disassembling some parts.

Moreover, there is not just one sequence: different evolutionary lineages evolving in parallel can produce different DCKs. According to the "Symbiogenesis" theory, different DCKs produced independently can sometimes merge to support new forms of life combining different evolutionary strands.²⁸ 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.

²⁸ <http://en.wikipedia.org/wiki/Symbiogenesis>

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.

5.2 Euclid's construction kit

An old example of great significance for science, mathematics, and philosophy, 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. This construction kit makes it possible to bisect, but not trisect, an arbitrary planar angle.

However, a slight extension to Euclidean geometry, the "Neusis construction", known to Archimedes, allows line segments to be translated and rotated in a plane while preserving their length, and certain incidence relations. This allows arbitrary angles to be trisected! (See <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html>)

The ability of humans to discover such things must depend on evolved information processing capabilities of brains that are as yet unknown and not yet replicated in AI reasoning systems. The idea of a space of possibilities generated by a physical construction kit may be easier for most people to understand than the comparison with generative powers of grammars, formal systems, or geometric constructions, though the two are related, since grammars and mathematical systems are abstract construction kits that can be parts of hybrid construction kits.

Concrete construction kits corresponding to grammars can be built out of physical structures: for example a collection of small squares with letters and punctuation marks, and some blanks, can be used to form sequences that correspond to the words in a lexicon. A cursive ("joined up") script requires a more complex physical construction kit. Human sign-languages are far more demanding, since they involve multiple body parts moving concurrently.

[Expand the following:]

Some challenges for construction kits used by evolution, and also challenges for artificial intelligence and philosophy, arise from the need to explain both (a) 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 (b) how various subsets of those organisms (e.g. nest-building birds) developed specific topological and geometrical *reasoning* abilities used in controlling actions or solving problems, and finally (c) 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*^[1].

There are important aspects of those mathematical competences that, as far as I know, have not yet been replicated in Artificial Intelligence or Robotics²⁹. Is it possible that currently understood forms of digital computation are inadequate for the tasks, whereas chemistry-based information processing systems used in brains are richer, because they combine both discrete and continuous operations, as discussed in Section 2.5? (That's not a rhetorical question: I don't know the answer.)

²⁹ Some of them listed in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/mathstuff.html>

5.3 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 finite region, 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. Kant had no access to a 20th century formal axiomatisation of Euclidean geometry. What he, and before him Euclid, Archimedes and others had were products of evolution. What products?

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

For example, in a video lecture that may be hard for non-mathematicians to follow, [Dana Scott, 2014](#) presents a series of constructions building up Euclidean geometry using a very different set of initial ideas (concepts and axioms) from those in Euclid's Elements. In doing so he uses an abstract derived construction kit that is not an ancient product of biological evolution, but a very recent product of human mathematical thinking (though making use of much older mathematical capabilities). One of its features is that it does not start from points and measurable distances, angles, areas, etc., but constructs them from more primitive (in that system) topological notions. It also seems to be capable of generating (a model of) the natural numbers and the continuum. This is a very late, distant, product of very much older construction kits produced more directly by evolution, though newer kits seem to require collaborative interactions between members of a mathematical community: i.e. they are in part products of social/cultural evolution, even though humans and cultures are not mentioned in them.

In order to understand biological evolution on this planet we need to understand the many sorts of construction kit 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). Somehow these also made possible the creation and use of more abstract construction kits, such as those used in production of human languages and methods of construction and exploration of mathematical domains.

An open research question is whether a construction kit capable of producing all the non-living structures that originally formed on the planet would also suffice for evolution of all the forms of life on this planet, or whether life and evolution have additional requirements, e.g. some aspect of cosmic radiation?

5.4 Evolution's (blind) mathematical discoveries

Some construction kits and their products have mathematical properties, so life and mathematics are closely connected. More complex relationships arise after evolution of mathematical metacognitive mechanisms. 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 occurrences. Proofs are often thought of as abstract entities that can be represented physically in different ways (using different formalisms) for communication, persuasion (including self-persuasion), predicting, explaining and planning. 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 proof that some construction kit makes that sort of entity possible. The evolutionary or developmental trail, like a geometric construction, answers the question: "how is that sort of thing possible?"

So biological evolution can be construed as a "blind theorem prover", despite there being no intention, or explicit recognition, regarding 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 [[Sloman 1962](#)]. This answers questions like "How is it possible for things that make mathematical discoveries to exist?", an example of explaining a possibility (See Note 4). Attempting to go too directly from hypothesized properties of the primordial construction kit to explaining advanced capabilities such as human self-awareness, without specifying all the relevant construction kits, including required temporary scaffolding (Sect. [2.7](#)) will fail, because short-cuts omit essential details of both the problems and the solutions, like mathematical proofs with gaps.

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" started happening long before there were any humans doing mathematics (which refutes claims that mathematics is a human creation). 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 discovery that something important is possible does not require collection of statistics: A single instance suffices. And no amount of statistical evidence can show that something is impossible: structural constraints need to be understood.

There is, however, purely mathematical notion of probability that is defined in terms of ratios of sizes of sets of possibilities, either discrete countable sets or measurable sets. An example of the first case is the probability of throwing a total number divisible by five with two dice. Another is the probability of getting three heads in succession in ten tosses of a coin. In these cases the probability can be calculated from the specification of the problem: no experiments are necessary.

For human evolution, a particularly important subtype of mathematical discovery was unwitting discovery and use of mathematical (e.g. topological) structures in the environment, a discovery process that starts in human children before they are aware of what they are doing, and in some species without any use of language for communication. Examples are discussed in the "Toddler Theorems" document [[Sloman 2013c](#)].

Dana Scott's new route to Euclidean geometry

(Added 3 Sep 2017)

In a video lecture parts of which may be hard for non-mathematicians to follow, [Dana Scott, 2014](#) presents a series of constructions building up Euclidean geometry using a very different set of initial ideas (concepts and axioms) from those in Euclid's Elements. In doing so he uses an abstract derived construction kit that is not an ancient product of biological evolution, but a very recent product of human mathematical thinking (though making use of much older mathematical capabilities). One of its features is that it does not start from points and measurable distances, angles, areas, etc., but constructs them from more primitive (in that system) topological notions concerned with spatial regions, their relationships, and ways of combining them. This is an example of mathematical discovery that does not start from some set of logically formulated axioms and rules of inference: instead it starts from ways of thinking about spatial structures and relationships, initially involving only spatial regions.

It also seems to be capable of generating (a model of) the natural numbers and the continuum. This is a very late, distant, product of very much older construction kits produced more directly by evolution, though newer kits seem to require collaborative interactions between members of a mathematical community: i.e. they are in part products of social/cultural evolution, even though humans and cultures are not mentioned in them.

[CONTENTS](#)

6 Varieties of Derived Construction Kit

DCKs may differ (a) at different evolutionary stages within a lineage, (b) across lineages (e.g. in different coexisting organisms), and (c) during development of an individual that starts as a single cell and develops mechanisms that support different kinds of growth, development and information processing, at different stages of development (Section [2.3](#)). New construction kits can also be produced by cultures or ecosystems (e.g. human languages) and applied sciences (e.g. bioengineering, computer systems engineering). New cases build on what was previously available. Sometimes separately evolved DCKs are combined, for instance in symbiosis, sexual reproduction, and individual creative learning.

What sort of kit makes it possible for a child to acquire competence in any one of the thousands of different human languages (spoken or signed) in the first few years of life? Children do not merely *learn* pre-existing languages: they *construct* languages that are new for them, constrained by the need to communicate with conspecifics, as shown dramatically by Nicaraguan deaf children who developed a sign language going beyond what their teachers understood [Senghas 2005]. There are also languages that might have developed but have not (yet). Evolution of human spoken language may have gone from purely internal languages needed for perception, intention, etc., through collaborative actions then signed communication, then spoken communication, as argued in [Sloman 2008/2011].

If language acquisition were mainly a matter of learning from expert users, human languages could not have existed, since initially there were no expert users to learn from, and learning could not get started. This argument applies to any competence thought to be based entirely on learning from experts, including mathematical expertise. So data-mining in samples of expert behaviours will never produce AI systems with human competences - only inferior subsets at best, though some narrowly focused machines based on very large data-sets or massive computational power may outperform humans (e.g. IBM's Deep Blue chess machine and WATSON).³⁰

³⁰ This comparison needs further discussion. See <http://www.popsci.com/science/article/2012-12/fyi-which-computer-smarter-watson-or-deep-blue>

The history of computing since the earliest calculators illustrates changes that can occur when new construction kits are developed. There were not only changes of size, speed and memory capacity: there have also been profound qualitative changes, in new layers of virtual machinery, such as new sorts of mutually interacting causal loops linking virtual machine control states with portions of external environments, as in use of GPS-based navigation. Long before that, evolved virtual machines provided semantic contents referring to non-physical structures and processes, e.g. mathematical problems, rules of games, and mental contents referring to possible future mental contents ("What will I see if...?") including contents of other minds.

I claim, but will not argue here, that some new machines cannot be *fully described* in the language of the FCK even though they are *fully implemented* in physical reality. (See Section 2.2 on ontologies.) We now understand many key components and many modes of composition that provide platforms on which *human-designed* layers of computation can be constructed, including subsystems closely but not rigidly coupled to the environment (e.g. a hand-held video camera).

Several different "basic" abstract construction kits have been proposed as sufficient for the forms of (discrete) computation required by mathematicians: namely Turing machines, Post's production systems, Church's Lambda Calculus, and several more, each capable of generating the others. The Church-Turing thesis claims that each is sufficient for all forms of computation.³¹ 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. One interpretation of the Church-Turing thesis is that these construction kits generate all *possible* forms of information processing.

³¹ For more on this see: http://en.wikipedia.org/wiki/Church-Turing_thesis

But it is not at all obvious that those discrete mechanisms suffice for all biological forms of information processing. For example, chemistry-based forms of computation include both discrete mechanisms (e.g. forming or releasing chemical bonds) of the sort Schrödinger discussed, and continuous process, e.g. folding and twisting used in reproduction. [Ganti 2003] shows how a chemical construction kit' can support forms of biological information processing that don't depend only on external energy sources (a feature shared with battery-powered computers), and also supports growth and reproduction using internal mechanisms, which human-made computers cannot do (yet).

I argued, in [Sloman \[IRREL\]](#), that Turing machines, as originally specified by Alan Turing, are irrelevant to AI, and that most of the ideas about computation construed as a form of information processing that have informed AI did not depend on the theory of Turing machines or the existence of physical Turing machines. Two outstanding examples of books written about biological information processing without any prior experience of working computers as we now understand them are [Fraud\(1943\)](#) and [\[Schrödinger 1944\]](#). I suspect Turing would have agreed with this, especially given the direction his thought seemed to be taking in 1952 and his remark in [\[Turing 1950\]](#) that *"In the nervous system chemical phenomena are at least as important as electrical."*

In organisms, there seem to be many different sorts of construction kit that allow different sorts of information processing 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.³²

³² Examples of human mathematical reasoning in geometry and topology that, as far as I know, still resist replication in computers are presented in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html>, and <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html>

6.1 A new type of research project

Many biological processes and associated materials and mechanisms are not 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 suffice to explain the possibility of all known evolved biological DCKs.

Researchers in fundamental physics or cosmology do not normally attempt to ensure that their theories explain the many materials and process types that have been explored by natural selection and its products, in addition to known facts about physics and chemistry.

[\[Schrödinger 1944\]](#) pointed out that a theory of the physical basis of life should explain how such phenomena are possible, though he could not have appreciated some of the requirements for sophisticated forms of information processing, because, at the time he wrote, scientists and engineers had not learnt what we now know.

(Curiously, although he mentioned the need to explain the occurrence of metamorphosis in organisms, the example he gave was the transformation from a tadpole to a frog. He could have mentioned more spectacular examples, such as the transformation from a caterpillar to a butterfly via an intermediate stage as a chemical soup in an outer case, from which the butterfly later

emerges.³³)

³³ <http://en.wikipedia.org/wiki/Pupa>
<http://en.wikipedia.org/wiki/Holometabolism>

[Penrose 1994] attempted to show how features of quantum physics explain obscure features of human consciousness, especially mathematical consciousness, but he ignored all the intermediate products of biological evolution on which animal mental functions build. Human mathematics, at least the ancient mathematics done before the advent of modern algebra and logic, seems to build on animal abilities, such as abilities to see various types of affordance. The use of diagrams and spatial models by Penrose may be an example of that.

It is unlikely that there are very abstract human mathematical abilities that somehow grow directly out of quantum mechanical aspects of the FCK, without depending on the mostly unknown layers of perceptual, learning, motivational, planning, and reasoning competences produced by billions of years of evolution.

20th century biologists understood 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.³⁴

³⁴ [http://en.wikipedia.org/wiki/J. B. S. 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 biological construction kits may continue indefinitely, producing new kits with components and modes of composition that allow production of increasingly complex types of *structure* and *behaviour* in organisms. That idea is familiar to computer scientists and engineers since thousands of new sorts of computational construction kit (new programming languages, new operating systems, new virtual machines, new development toolkits) have been developed from old ones in the last half century, making possible new kinds of computing system that could not previously be built from the original computing machinery without introducing new intermediate layers, including 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. [Sloman 2013a] discusses the importance of virtual machinery in extending what information processing systems can do, and the properties they can have, including radical self-modification while running.

6.2 Construction-kits for biological information processing

Each newly evolved DCK enables evolution of new physical/chemical mechanisms that need to be controlled. This provides opportunities for evolution to produce additional mechanisms for information-based control of those mechanisms, for example choosing between competences or "tuning" them dynamically, often on the basis of their mathematical properties. Almost any new development can provide opportunities for evolution to produce yet another level of information-based control --- a recurring process.

In some organisms, implicit mathematical discovery processes enable production of competences used in new kinds of generic understanding of sensory information, e.g. locating perceived objects and events in space, and recognition of possibilities for satisfying needs or avoiding damage by changing locations or relationships, or selecting trajectories.

Later developments might support synthesis of separate information fragments into coherent wholes. This would provide new functionality for control systems generating motives, constructing plans, selecting and controlling behaviour, and making predictions. These apparently simple descriptions could cover many very complex changes over thousands or millions of generations, many of which might diverge into different evolutionary trajectories, using different physical forms and resources, and using different sorts of information and different ways of processing information.

I conjecture that during such processes evolution made many mathematical "discoveries" that were "compiled" into designs producing useful behaviours, e.g. use of negative feedback loops controlling temperature, osmotic pressure and other states, use of geometric constraints by bees whose cooperative behaviours produce hexagonal cells in honeycombs, and enriching ontologies for distinguishing situations requiring different behaviours, e.g. manipulating different materials, hunting different kinds of prey, or coping with new sorts of terrain.

Later still (much later!), new sorts of evolved construction kit produced metacognitive mechanisms enabling individuals to notice and reflect on their own discoveries (enabling some of them to notice and remove flaws in their reasoning).

In some cases those metacognitive capabilities allowed individuals to communicate discoveries to others, discuss them, and organise them into complex highly structured bodies of shared knowledge, such as Euclid's *Elements* (Note [1](#)).

I don't think anyone knows how long all of this took, what the detailed evolutionary changes were, or how the required mechanisms of perception, motivation, intention formation, reasoning, planning and retrospective reflection evolved. 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.

That includes seeking unnoticed overlaps between the human competences discovered by metacognitive mechanisms, and similar competences in animals that lack the metacognition. Examples include young humans making and using mathematical discoveries, on which they are unable to reflect because the required metacognitive architecture has not yet developed, and similar discoveries in other intelligent species that seem able to make and use similar "proto-mathematical" discoveries, without the meta-cognitive abilities required to notice what they are doing, such as squirrels, crows, elephants and apes. These ideas could stimulate new research in robotics attempting to replicate such competences. [REF Richard Byrne]

Most of these naturally occurring (proto-)mathematical abilities have not yet been replicated in Artificial Intelligence systems or robots, unlike logical, arithmetical, and algebraic competences that are relatively new to humans and (paradoxically?) easier to replicate on computers.

Examples of topological reasoning (e.g. about equivalence classes of closed curves) that are not yet modelled in computers (as far as I know) are referenced in Note [32](#).

Even the ability to reason about alternative ways of putting a shirt on a child, and their implications (Note [18](#)), is still lacking.

It is not clear whether the difficulty of replicating such mathematical reasoning processes is 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 computers to replicate some of the products of biological evolution, or both!

Perhaps there are important forms of representation or types of information processing architecture still waiting to be discovered by AI researchers and software engineers.

Alternatively, the gaps may be connected with properties of chemistry-based information processing mechanisms combining discrete and continuous interactions, or other physical properties that cannot be replicated exactly (or even approximately) in familiar forms of computation. (This topic requires more detailed mathematical analysis.)

6.3 Representational blind spots of many scientists

Although I 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. Even the pioneering work in [\[Schrödinger 1944\]](#) focused only on a tiny subset of information processing requirements concerned with reproduction and control of individual development.

Presentations on dynamics of physical systems generally make deep use of branches of mathematics concerned with numerical values, and the ways in which different measurable or hypothesized physical values, or their derivatives, or in various combinations, do or do not co-vary, as expressed in (probabilistic or non-probabilistic) equations of various sorts. But the biological functions of complex physiological structures, especially structures that change in complexity, and structures concerned with processing many varieties of information, don't necessarily have those forms.

For example they may make use of logical or algebraic structures and processes, or topological structures and processes such as hole formation or formation of chains made of linked, or threaded components (as in spinal chords). How to get computer-based AI systems to reason about such things in human-like ways is a difficult challenge.

Biological mechanisms that may involve structure-formation and manipulation 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.

Forms of mathematics used by physicists are not necessarily useful for studying such 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 those aspects of biology, and therefore for specifying what physicists need to explain.

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 well supported by widely used tools such as Matlab, but is not necessarily best suited for the majority of types of mental content (including grammatical and semantic structures of thoughts, intentions, questions, etc.), 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 that are now being extended in computer models of changing, interacting molecular structures. To see examples, search for online videos of computer simulations of chemical reactions including protein folding processes.

6.4 Representing rewards, preferences, values

It is often assumed that all intelligent decision making uses positive or negative scalar rewards or utility values that are comparable across options [[Luce & Raiffa 1957](#)]. But careful attention to consumer magazines, political debates, and the varieties of indecision that face humans throughout 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. Analysis of examples in consumer reports led to the conclusion that "better" is a complex, polymorphic, logical concept with a rich structure that cannot be reduced to comparison of numerical values [[Sloman 1969](#),[Sloman 1970](#)]. Instead of a linear reward or utility metric, choices for intelligent individuals, or for natural selection, involve a complex network of partial orderings, with "annotated" links between nodes (e.g. "better" qualified by conditions: "better for", "better if", "better in respect of", ...). In the Birmingham Cognition and Affect project [[Sloman 2003](#)], those ideas informed computational models of simple agents with complex choices to be made under varying conditions, but the project merely scratched the surface, as reported in [[Beaudoin Sloman 1993](#), [Beaudoin 1994](#), [Wright, Sloman, Beaudoin. 1996](#), [Wright 1997](#)]. Most AI/Cognitive Science models use much shallower notions of motivation.

Despite all the sophistication of modern psychology and neuroscience, I believe they currently lack the conceptual resources required to describe either *functions* of brains in dealing with these matters, including forms of development and learning required, or the *mechanisms* implementing those functions. In particular, we lack deep explanatory theories about mechanisms that led to mathematical discoveries over thousands of years, including brain mechanisms producing mathematical conjectures, proofs, counter-examples, proof-revisions, new scientific theories, new works of art and new styles of art. In part that's because models considered so far lack both sufficiently rich forms of information processing (computation), and sufficiently deep methodologies for identifying what needs to be explained. There are other unexplained phenomena concerned with artistic creation and enjoyment, mechanisms involved in finding something funny, and mechanisms by which internal conflicts may continue after decisions have been taken, e.g. when rejected desires continue to grab attention. But those will not be pursued here.

[CONTENTS](#)

7 Computational/Information-processing construction kits

Since the mid 20th century we have been learning 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 crucial, but poorly understood, feature of that history is the growing use of construction kits based on virtual machinery, mentioned [previously](#).

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 could be used by evolution to produce an increasing variety of types of *virtual* machinery as well as increasingly varied *physical* structures and mechanisms.

7.1 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 is essential for life as we know it, and also for forms of information processing produced by evolution (mostly not yet charted).

Research in fundamental physics is a search for the construction kit that has the generative power to accommodate all the possible forms of matter, structure, process and 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 study structures and processes in living matter, or processes of evolution, development, perception, reasoning, learning and mathematical discovery found in living organisms. Most physicists seem not to be interested in trying to ensure that their fundamental theories can account for those features of living matter also. There are notable exceptions, such as Schrödinger, Penrose, Deutsch, Stegmark and others[REFS], but most physicists who discuss physics and life (in my experience) understandably ignore 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 and change, the ways in which products of evolution become more complex and more diverse over time, and the many kinds of information processing and control in individuals, in colonies (e.g. ant colonies), societies, and ecosystems.

If cosmologists and other theoretical physicists attempted to account for a wide range of biological phenomena (including the phenomena discussed here in connection with the Meta-Morphogenesis project) I suspect that 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 physical universe (or the Fundamental Construct Kit). In part that could be a consequence of the forms of mathematics known to physicists. The well known challenge by Nobel Prize winning physicist Philip Warren Anderson [["More is different"\(1972\)](#)], discussed briefly [below](#), is also relevant.

It may take many years of research to find out what exactly is missing from current physical theories. 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 unending attempts to put computation to new uses. Collecting detailed examples of such challenges helps us assemble tests to be passed by future proposed theories: a catalogue of realised and conjectured possibilities that a deep physical theory needs to be able to explain. Examples I have been collecting include types of mathematical discovery and reasoning that are not easily implementable using current AI techniques. Here are two examples:

A pre-verbal toddler with a pencil explores 3D topology.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#pencil>

A rubber band impossibility

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/rubber-bands.html>

Understanding the evolution of such capabilities and the information processing mechanisms they require is a hard challenge:

http://www.cs.bham.ac.uk/research/projects/cogaff/sloman_iccm_17.pdf

Many of the people who think that computers are *ideally* suited to replicating human mathematical abilities, because they have a narrow view of mathematics, including perhaps arithmetic, algebra and logic, and do not recognize the non-numerical mathematical competences present even in pre-verbal children (some of them studied by Piaget and his students, e.g. in Piaget [[1981](#), [1983](#)]). Compare the document on "Toddler Theorems" ([Note 24](#)).

Perhaps the most tendentious proposal here is that an expanded physical theory, instead of being expressed mainly in terms of equations relating numerical measurables of various kinds, will 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, such as possible future forms of intelligence. (I suspect that theoretical computer science will be one of the sources of relevant new forms of non-logical, non-algebraic, representation.)

Theories of motivation based on utility measures may need to be enhanced, or replaced with new theories of how benefits, evaluations, comparisons and preferences, can be expressed [[Sloman 1969](#)]. We must also avoid assuming optimality -- a common trap. Evolution produces designs as diverse as microbes, cockroaches, elephants and orchids, none of which is optimal or rational in any simple sense, yet many of them survive and sometimes proliferate, not because they are optimal, but because they are lucky, at least for a while. The same applies to human decisions, policies, preferences, cultures, gambles, conjectures, etc.

8 Types and levels of explanation of possibilities

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 a certain sort of crane 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 and economy of specification.

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: it is merely an experimental foray, preparing the ground for future work:

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 X (e.g. 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 this with 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 piecewise relations, where the whole structure cannot possibly exist, because the relations cannot hold simultaneously, e.g. X is above Y, Y is above Z, Z is above X. It is possible to depict such objects, e.g. in pictures of impossible objects by Reutersvard, Escher, Penrose, and others.^{[35](#)}

^{[35](#)} See [\[Sloman 2015\]](#)

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 with restricted practical uses, e.g. strongly typed programming languages. For some purposes less restricted languages, needing greater care in use, are preferable, including human languages, as I tried to show in [\[1971\]](#).

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 scaffolding (Sect. [2.7](#)) 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. There may be several different routes to the same end result.

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 are 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 each state-pair may be shared by a vast, or even infinite, collection of different possible trajectories producing the same end state, e.g. in a continuous space.

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 several variables. If trajectories include passing through holes, or 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, for example, the different ways of threading a shoe lace with the same end result.

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, or at least searching for, *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, possibly during execution of a purely qualitative plan.

These remarks are closely related to the origins of abilities to reason about geometry and topology.³⁶

³⁶ Illustrated in these discussion notes:

<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 "structural conformity" and "process possibility", discussed above. 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. For example, 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 15, that will collapse 1,307,674,368,000 different sequences into one.

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. In more complex cases, the end state may depend on the relative speeds of sub-processes and also continuously changing spatial relationships. 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.

For some construction kits, certain constructs made of a collection of sub-assemblies may require different sub-assemblies to be constructed in parallel, for example if completing some too soon would make the required final configuration unachievable, like completing rings separately before using them to form a chain.

5 Iterative or recursive abstraction: Some process types involve unspecified numbers of parts or 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 a large variety of possibilities inherent in a construction kit' in an extremely powerful and economical way. Many new forms of abstraction of this type have been discovered by computer scientists developing programming languages, for operating not only on numbers but many other structures, e.g. trees and graphs.

Evolution may also have "discovered" many cases, 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, or a few discoveries, at a high enough level of abstraction.

As mentioned previously, some related abstractions may have resulted from parametrisation: processes by which details are 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 also Section [2.3](#) on epigenesis.)

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 required steps or required end states, 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).

7 Use of explicit 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).

It seems that some of the reproductive mechanisms that depend on parental care make use of mechanisms that generate intentions and possibly also plans in carers, for instance intentions to bring food to an infant, intentions to build nests, intentions to carry an infant to a new nest, intention to migrate to another continent when temperature drops, and many more. Use of intentions that can be carried out in multiple ways selected according to circumstances rather than automatically triggered reflexes could cover a far wider variety of cases, but would require provision of greater intelligence in individuals.

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.

Concluding comment on explanations of possibilities:

Those are all examples of components of explanations of assembly processes, including self-assembly. 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 human development, and to a lesser extent development in other species, there are many additional sub-cases involving construction kits both for creating information structures and creating information processing mechanisms of many kinds, including perception, learning, motive formation, motive comparison, intention formation, plan construction, plan execution, language use, 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 all 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, especially when followed by caring for offspring. The latter allows a parent's influence to continue during development, as does teaching of younger individuals by older ones. This also allows development of cultures suited to different environments.

To conclude this rather messy section: the investigation of different types of generality in modes of explanation for possibilities supported by a construction kit is also relevant to modes of specification of new designs based on the kit. Finding economical forms of abstraction may have many benefits, including both reducing search spaces when trying to find a new design and also providing a generic design that covers a broad range of applications tailored to detailed requirements. Of particular relevance in a biological context is the need for designs that can be adjusted over time, e.g. during growth of an organism, or shared across species with slightly different physical features or environments. Many of the points made here are also related to structural changes in types of computer programming language and software design specification language. Evolution may have beaten us to important ideas. That these levels of abstraction are possible is a metaphysical feature of the universe, implied by the generality of the FCK.

[CONTENTS](#)

9 Alan Turing's Construction kits

[[Turing 1936](#)] showed 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 into 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,^{[37](#)} a construction kit that creates changing patterns in 2D regular arrays.

³⁷ <http://en.wikipedia.org/wiki/Conway.27s.Game.of.Life>

Stephen Wolfram has written a great deal about the diversity of constructions that can be explored using such cellular automata. 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 able to 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 would change.

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 (e.g. from keyboards, buttons, joysticks, mice, electronic piano keyboards, network connections, and many more) and can also send signals to external devices. Theorems about disconnected Turing machines may not apply to machines with rich two-way interfaces connected to external environment.

Turing machines and Game of Life machines can be described as "self-propelling" because once set up they can be left to 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 physical machines with an external power supply. In contrast, [\[Ganti 2003\]](#) 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 that soar. 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.

9.1 Beyond Turing machines: chemistry

Turing also explored other sorts of construction kit, including types of neural net and extended versions of Turing machines with "oracles" added. Shortly before his death (in 1954), he published [\[Turing 1952\]](#) 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. I have been trying to show how that can be seen as a very simple example of something far more general.

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 many who have played with construction kits will have noticed - e.g. when a partial construction produces a gap whose width does not exactly match the width of any available pieces).

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".³⁸

³⁸ One of many online explanations is
<http://www.theprojectspot.com/tutorial-post/simulated-annealing-algorithm-for-beginners/6>

A recently published book [Wagner 2014] claims 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 being 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.³⁹

³⁹ An interview with the author is online at
<https://www.youtube.com/watch?v=wyQgCMZdv6E>

I have not yet checked whether the book discusses uses of abstraction and the evolution of mathematical and meta-mathematical competences discussed here. Nevertheless, it seems to be an (unwitting) contribution to the Meta-Morphogenesis project.

Recent work by Jeremy England at MIT, reported by Natalie Wolchover⁴⁰, may turn out also to be relevant, by extending the ideas in *What is life?* [Schrödinger 1944] by using Quantum theory to explain how it is possible for some important precursors of life to come into existence on a lifeless planet.

⁴⁰ Natalie Wolchover, 2014, Report on work by Jeremy England, "This Physicist Has A Groundbreaking Idea About Why Life Exists" in *Business Insider* Dec. 8, 2014,
<http://www.businessinsider.com/groundbreaking-idea-of-lifes-origin-2014-12>
<https://www.quantamagazine.org/20140122-a-new-physics-theory-of-life/>

Some of the structures that might spontaneously form could be building blocks not only for some of the earliest forms of life (as described by Ganti (1971/2003)) but possibly also for some of the "construction kits" and forms of scaffolding required for biological evolution. (My thanks to Aviv Keren for drawing my attention to the paper on England.)

9.2 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 including proofs that end in important 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 plasticine, or soap-film.

9.3 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, whether certain pairs of numbers can be adjacent, or whether the same number can occur in more than one square. For a given grid and a given set of constraints, there 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 kit have been investigated by mathematicians, logicians, linguists, computer scientists, musicians and other artists.

Analysis of chemistry-based construction kits for information processing systems would range over a far larger class of possible systems than Turing machines (or digital computers), 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]. I don't know whether anyone has a deep theory of the scope and limits of chemistry-based information processing.

Recent discoveries indicate that some biological mechanisms use quantum-mechanical features of the FCK that we do not yet fully understand, providing forms of information processing that are very different from what current computers do. E.g. a presentation by Seth Lloyd, summarises quantum phenomena used in deep sea photosynthesis, avian navigation, and odour classification.⁴¹ This may turn out to be the tip of an iceberg of quantum-based information processing mechanisms.

⁴¹ <https://www.youtube.com/watch?v=wcXSpXyZVuY>

There are some unsolved, very hard, partly ill-defined, problems about the variety of functions of biological vision: e.g. simultaneously interpreting a very large, varied and changing collection of visual fragments, perceived from constantly varying viewpoints, e.g. as you walk through a garden with many unfamiliar flowers, shrubs, bushes, etc. moving irregularly in a changing breeze. Could some combination of quantum entanglement and non-local interaction play a role in rapidly and simultaneously processing a large collection of mutual constraints between multiple visual fragments? The ideas are not yet ready for publication, but work in progress is recorded here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/quantum-evolution.html>.

Some related questions about perception of videos of fairly complex moving plant structures are raised here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision/plants/>.

9.4 Towers vs forests (Draft: added 2 Feb 2016)

[Skip on first reading?]

As more and more DCKs evolve (or are developed in individuals) they can be thought of as forming

towers growing "upwards" from the Fundamental construction kit. However, towers, like most naturally formed non-living structures that grow upward, tend to become narrower the higher they grow, whereas construction kits become larger and more complex, as required for building larger and more complex organisms, including their information processing mechanisms, as well as increasingly many physical mechanisms and functions. Moreover processes of differentiation during evolution will cause towers of construction kits to form branches and proliferate as they grow up, as trees and bushes do.

In the history of computing there are two directions of growth of towers of abstraction, described in [\[Fisher,Piterman,Vardi, 2011\]](#)(FPV). The growth of *engineering know how* in a community moves from small detailed mechanisms for performing simple tasks (e.g. simple arithmetic operations, or sorting operations) through designs of increasing complexity adding higher levels that specify more powerful functionality built on previously understood lower levels. This is a kind of *upward* growth.

In contrast the production of *a new working system* using existing general know-how often starts from a specification of required high-level functionality, moving "downwards" through layers of increasingly complex detail, eventually specifying actual physical components and their physical relationships -- a process of *downward* growth. (In many cases discoveries made during "downward design" and implementation cause higher level goals to be re-formulated, and/or construction kits to be extended.)

However, the bottom layers generally use reconfigurable off-the-shelf hardware systems that can be configured for particular applications by setting switches -- a process now done automatically by loading a program into memory. But, as [\[FPV\]](#) point out, sometimes new special purpose physical machinery has to be assembled to provide the bottom layer, in some cases using a mixture of digital and analog (continuously variable) machinery.

These upward and downward processes of design and construction can be used for designing and building a specific type of machine. However that is rarely done by individual users, because a by-product of such processes is *re-usable* knowledge about kinds of functionality and how to achieve them, leading to production of off-the-shelf physical components and off-the-shelf software components, at different levels of abstraction: compilers, interpreters, operating systems, networking hardware and software, configuration tools, product development toolkits, testing frameworks (including benchmarking data), and many "pre-packaged" applications that can be acquired and used as components (often without even understanding exactly what they do or how they do it).

Increasingly these pre-packaging processes have been applied to *virtual* machinery, or useful components of virtual machinery. The virtual machines themselves are not stored or transferred when acquired. Rather computer programs ("code"), usually in the form of textual specifications encoded in bit patterns, are stored in one or more "central" repositories and then downloaded to a variety of physical machines (computers, tablets, smart phones, and other devices) where the VMs can be set up and run. The transported code need not (and most often will not) consist of machine instructions for the target device. Rather some relatively more abstract programming language is used for the new extensions ("apps").

This works if the receiving device already has facilities (e.g. appropriate virtual machinery) for setting up and running virtual machines specified in the VM language. Typically such a newly installed package will need to use pre-existing hardware and software interfaces to devices like

screen, keyboard, touchpad, camera, microphone, and to other (virtual machine) packages, e.g. virtual keyboard, memory management system, audio and video packages, internet connection packages, updating mechanisms, security checks, etc.

I suspect that all of these recent advances make use of special cases of types of design and mechanisms for combining designs, that biological evolution produced and used (with many variations) millions of years ago. Without that packaging and reuse of powerful abstract design specifications, it would have been impossible to produce the many amazing evolutionary lineages found in fossil records -- e.g. the extraordinary variety of highly sophisticated flying mechanisms in insects, some of them illustrated here.

<https://www.youtube.com/watch?v=LCekEHI82oA>

Animal Flight: Wing Structures and Wingbeat Mechanisms by Larry Keeley.

<https://www.youtube.com/watch?v=ARLTxG2gh3I>

Life on Earth - e03 - Flight control in flies

What is easy to forget is that the use of those physical mechanisms require control "software" to allow the variety of behaviours based on flying (including catching prey, escaping from predators, and even mid-air mating in some species). An example of a control problem is detecting something dangerous approaching from a certain direction, choosing an escape direction, and triggering actions rapidly causing motion in that direction.

Somehow the physical mechanisms and the control packages must have evolved together. A possible partial explanation of how that happened may be that evolution "discovered" types of decomposition of designs into abstract re-usable features and parameters that specify particular uses of those features. The physical design parameters and the control parameters would have to vary together as organisms evolved, while preserving important control relationships.

I suspect evolution discovered forms of re-usable design specification for these purposes that human scientists and engineers have not yet re-discovered, though they have discovered similar (probably much simpler?) relationships between physical designs and control mechanisms, for use in families of aircraft, e.g. fighter planes required for different types of combat, and helicopters of different sizes and functions. I am not an expert in any of these domains: I am merely reasoning from fairly abstract general principles relating control functions and physical structures, as illustrated, for example, here:

<http://www.helicoptermaintenancemagazine.com/article/understanding-helicopter-automatic-flight-control-systems-afcs>

Understanding Helicopter Automatic Flight Control Systems (AFCS)

Compare the control problems facing an octopus:

<http://nymag.com/scienceofus/2017/03/whats-it-like-to-be-an-octopus.html>

(NOTE (Aug 2017) Technological advances have made possible a huge variety of recordings of different flight patterns in insects and other organisms, and investigations of the physical mechanisms involved. I suspect that the control mechanisms used will turn out to be far more complex than motion control in apparently more intelligent animals such as apes, nesting birds, elephants, etc.)

In an advanced industrial society the many kinds of components used for building each functioning computing system will be designed, manufactured, and supplied by a host of different organisations, often spread across continents. In contrast as a biological organism grows, all the

initial materials, design specifications and practical construction know-how must be available in one place, though a developing organism may, over time, acquire new materials from different parts of its environment, possibly very widely dispersed (e.g. in migrating land mammals, whales and many bird species). Moreover, as pointed out in [\[Schrödinger 1944\]](#), many crucial materials with complex physical/chemical structures available in the environment may have been assembled at other times and places, including microorganisms on another part of the planet where the environment is very different.

This contrasts with a tower of abstraction that starts with a low level structure that has a lot of detail, and then moves "upwards" by forming more and more abstract and compact abstractions by replacing complex and detailed sub-structures with abstract specifications including variables that can be instantiated to produce multiple different instances.

[\[FPV\]](#) say: "*In our view, Biology needs to go beyond mere abstraction and develop its own tower of abstractions.*" I am in complete agreement though I would add to this "... and develop its own upwardly and downwardly branching towers of abstraction *for construction kits and for designs for organisms*". I think this is consistent with their intentions.

[CONTENTS](#)

10 Conclusion: Construction kits for Meta-Morphogenesis

A useful two-part survey by E.F.Keller [\[Keller 2008, Keller 2009\]](#), of previous attempts to show how life and its products relate to the physical world, concluded that attempts so far have not been successful. 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*".

I am trying to describe a long term, very ambitious, project to do that and more: including giving an account of branching layers of new *derived* construction kits produced by evolution, development and other processes. The physical world clearly provides a very powerful (physics-and-chemistry-based) 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.

The idea of a construction kit is offered as a new unifying concept for philosophy of mathematics, philosophy of science, philosophy of biology, philosophy of mind and metaphysics. One of the main aims is to explain how it is possible for minds to exist in a material world and to be produced by natural selection and its products. Related questions arise about the nature of mathematics and its role in life. The ideas are 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 by human engineers since around 1950 provide examples of non-reductive supervenience (as explained in [\[Sloman 2013a\]](#)). They are also useful

as relatively simple examples to be compared with far more complex products of evolution.

In [[Esfeld, Lazarovici, Lam, Hubert, 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. The ideas I have presented support a "multi-layer" version of the modal realist option (developing ideas in [[Sloman 1962](#), [Sloman 1978\(Ch.2\)](#), [Sloman 1996a](#), and [Sloman 2013a](#)]).

I suspect 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](#)], mentioned above, namely How should we 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".

However, the structural levels are not merely in the concepts used by scientists, but actually in the world. In the last half century or so we have learnt, for the first time, how to build complex functioning virtual machinery for which the same comment is relevant: [the study of such machinery is not applied electronics, or applied physics](#). See [[Sloman 2013a](#)] for more details.

We still have much to learn about the powers of the fundamental construction kit (FCK), including: (1) the details of how those powers came to be used for life on earth, (2) which sorts of derived construction kit (DCK) were required in order to make more complex life forms possible, (3) 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, (4) what possibilities the FCK has that have not yet been realised, (5) whether and how some version of the FCK could be used to extend the intelligence of current robots, (6) 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 (7) 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?^{[42](#)}

⁴² See the cartoon teasing particle physicists [above](#).

11 End Note (Turing's letter to Ashby)

While finishing off an earlier version of this paper I came across a letter Turing wrote to W. Ross Ashby in 1946 urging Ashby to use Turing's ACE computer to implement his ideas about modelling brains. Turing expressed a view that seems to be unfashionable among AI researchers at present (2015), but accords with the aims of this paper:

"In working on the ACE I am more interested in the possibility of producing models of the actions of the brain than in the practical applications to computing."

<https://www.bl.uk/collection-items/letter-from-alan-turing-to-w-ross-ashby>

(Thanks to Rodney Brooks for that link. I previously linked to this site, now defunct:

<http://www.rossashby.info/letters/turing.html>

It would be very interesting to know whether he had ever considered the question whether digital computers might be incapable of accurately modelling brains making deep use of chemical processes. He also wrote in [Turing 1950]

"In the nervous system chemical phenomena are at least as important as electrical."

But he did not elaborate on the implications of that claim.⁴³

⁴³ I think the ideas about "making possible" used here are closely related to Alastair Wilson's ideas about grounding as "metaphysical causation" [Wilson 2015].

12 Note on Barry Cooper

The Meta-Morphogenesis project, including this paper, owes a great debt to Barry Cooper who unfortunately died in October 2015. I had never met him until we both contributed chapters to a book published in 2011 on Information and Computation. Barry and I first met, by email, when we reviewed each others' chapters. Later, out of the blue, he invited me to contribute to the Turing centenary volume he was co-editing [Cooper Leeuwen 2013]. I contributed three papers. He then asked me for a contribution to Part 4 (on Emergence and Morphogenesis) based on Turing's paper on morphogenesis published in 1952, two years before he died. That got me wondering what Turing might have done if he had lived another 30-40 years. So I offered Barry a paper proposing "The Meta-Morphogenesis Project" as an answer. He accepted it (as the final commentary paper in the book) and ever since then I have been working full-time on the project. He later encouraged me further by inviting me to give talks and to contribute a chapter to this book. As a result we had several very enjoyable conversations. He changed my life, by giving me a new research direction, which does not often happen to 75-year old retired academics! (Now even older.) I wish we could continue our conversations.

Note added May 2017 My ideas have probably been influenced in more ways than I recognise by my colleague Margaret Boden, whose work has linked AI/Cognitive Science to Biology since the 1960s, notably in her extraordinary two volume survey published in 2006 by OUP, *Mind As Machine: A history of Cognitive Science*.

Ongoing work This paper is a snapshot (mid 2017) of a large and complex project, that continues to grow. Examples of continuing questioning, conjecturing, exploring, etc. can be found in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/maths-multiple-foundations.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/anger.html>

I also owe much to the highly intelligent squirrels and magpies in our garden, who have humbled me.

CONTENTS

FOOTNOTES:

(Footnotes now all relocated into the text)

References

[Anderson 1972]

Anderson P (1972) More is different. *Science*, New Series 177(4047):393-396,
http://robotics.cs.tamu.edu/dshell/cs689/papers/anderson72more_is_different.pdf

[Ashby 1952]

Ashby WR (1952), *Design for a Brain*, Chapman and Hall, London

[Ashby 1956]

Ashby WR (1956) *An Introduction to Cybernetics*. Chapman & Hall,
<http://pespmc1.vub.ac.be/ASHBBOOK.html>

[Bakermans (2015)]

Corien Bakermans, 2015, Chapter 1: Extreme environments as model systems for the study of microbial evolution, in *Microbial Evolution under Extreme Conditions*, Ed. Corien Bakermans. DeGruyter pp 153-170

[Beaudoin 1994]

Beaudoin L (1994) *Goal processing in autonomous agents*. PhD thesis, School of Computer Science, The University of Birmingham, Birmingham, UK,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#38>

[Beaudoin and Sloman 1993]

Beaudoin L, 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#16>

[Bell 2008]

Bell G (2008) *Selection The Mechanism of Evolution*. OUP, second Edition

[Bennett 2011]

Bennett K (2011) Construction Zone: No Hard Hat Required. *Philosophical Studies* 154:79-104,
DOI 10.1007/s11098-011-9703-8

[Chappell and Sloman 2007]

Chappell J, Sloman A (2007) Natural and artificial meta-configured altricial information processing systems. *International Journal of Unconventional Computing* 3(3):211-239,
<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717>

- [Chomsky 1965]
Chomsky N (1965) *Aspects of the theory of syntax*. MIT Press, Cambridge, MA
- [Coates et al 2014]
Coates J, Umm-E-Aiman, Charrier B (2014) Understanding "green" multicellularity: do seaweeds hold the key? *Frontiers in Plant Science*, Doi: 10.3389/fpls.2014.00737
- [Cooper/van Leeuwen(2013)]
S. B. Cooper and J. van Leeuwen, (Eds) (2013)
Alan Turing - His Work and Impact, Elsevier, Amsterdam.
Full contents, and further information, available here:
<http://www.cs.bham.ac.uk/~axs/amtbook/>
- [Cooper and Soskova 2017]
S.B. Cooper and M.I. Soskova (Eds), 2017, *The Incomputable: Journeys Beyond the Turing Barrier*
Springer-Verlag,
<http://www.springer.com/gb/book/9783319436678>
- Craik(1943)
Kenneth Craik, 1943, *The Nature of Explanation*, Cambridge University Press, London, New York,
- [Deacon 2011]
Deacon TW (2011) *Incomplete Nature: How Mind Emerged from Matter*. W. W. Norton & Company
- [Dennett 1995]
D.C. Dennett (1995) *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, Penguin Press, London and New York,
- [Esfeld et al in press]
Esfeld M, Lazarovici D, Lam V, Hubert M (in press) The Physics and Metaphysics of Primitive Stuff. *British Journal for the Philosophy of Science*
<http://arxiv.org/abs/1411.7545>
- [Fernando 2008]
Fernando C (2008) Review of "*The Principles of Life*" by Tibor Ganti. (2003, Oxford University Press.). *Artificial Life* 14(4):467-470, DOI 10.1162/artl.2008.14.4.14404
- [Fisher,Piterman,Vardi, 2011]
Jasmin Fisher and Nir Piterman and Moshe Y. Vardi, 2011,
The Only Way is Up: On A Tower of Abstractions for Biology, in
Proc. 17th International Conference on Formal Methods,
Eds. M. Butler and W. Schulte, pp. 3--11,
Springer-Verlag Berlin, Heidelberg, 978-3-642-21436-3,
<https://www.cs.rice.edu/~vardi/papers/fm11a.pdf>
- [Froese et al.(2014)]
T. Froese, N. Virgo, and T. Ikegami, (2014), Motility at the origin of life: Its characterization and a model', *Artificial Life*, 20(1), 55-76.
- [Ganti 2003]
Ganti T (2003), *The Principles of Life*. OUP, New York, Eds. Eörs Szathmáry & James Griesemer, Translation of the 1971 Hungarian edition
- [Goodwin 1995]
Brian Goodwin, 1995, Biology Is Just a Dance, in *The Third Culture: Beyond the Scientific Revolution*, Ed. John Brockman, Simon & Schuster, online at:
<https://www.edge.org/documents/ThirdCulture/k-Ch.4.html>

- [Gibson 1966]
Gibson J.J. (1966), *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston
- [Gibson 1979]
Gibson J.J. (1979) *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA
- [Kanefff 1970]
S. Kanefff, editor, *Picture language machines*, 1970, Academic Press, New York,
- [Kant 1781]
Kant I. (1781) *Critique of Pure Reason*. Macmillan, London, translated (1929) by Norman Kemp Smith
- [Karmiloff-Smith 1992]
Karmiloff-Smith A (1992) *Beyond Modularity: A Developmental Perspective on Cognitive Science*. MIT Press, Cambridge, MA
An informal, personal, discussion/review of this work is available here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html>
- [Kauffman 1993]
Kauffman, S (1993) *The Origins of Order: Self-Organization and Selection in Evolution* Oxford University Press
- [Kauffman 1995]
Kauffman S (1995) *At home in the universe: The search for laws of complexity*. Penguin Books, London
- [Keller 2008]
Keller EF (2008) Organisms, Machines, and Thunderstorms: A History of Self-Organization, Part One. in *Historical Studies in the Natural Sciences*, 38(1 (Winter)):45-75,
<http://www.jstor.org/stable/10.1525/hsns.2008.38.1.45>
- [Keller 2009]
Keller EF (2009) Organisms, Machines, and Thunderstorms: A History of Self-Organization, Part Two. Complexity, Emergence, and Stable Attractors. in *Historical Studies in the Natural Sciences* 39(1 (Winter)):1-31,
<http://www.jstor.org/stable/10.1525/hsns.2009.39.1.1>
- Kirschner & Gerhart(2005) M.W. Kirschner and J.C. Gerhart (2005), *The Plausibility of Life: Resolving Darwin's Dilemma*. Yale University Press: Princeton
- [Korthof 2003]
Korthof G (2003) Review of *The Principles of Life* by Tibor Ganti.
[updated 6 Oct 2014 http://wasdarwinwrong.com/korthof66.htm](http://wasdarwinwrong.com/korthof66.htm), updated 6 Oct 2014
- [Laird et al 1987]
Laird J, Newell A, Rosenbloom P (1987) SOAR: An architecture for general intelligence. *Artificial Intelligence* 33:1-64
- [Lakatos 1980]
Lakatos I (1980) Falsification and the methodology of scientific research programmes. In: Worrall J, Currie G (eds) *Philosophical papers, Vol I*, Cambridge University Press, Cambridge, pp 8-101
- [Luce and Raiffa 1957]
Luce R, Raiffa H (1957), *Games and Decisions: Introduction and Critical Survey* John Wiley and Sons, Inc.; Chapman and Hall, New York; London
- [Mathis et al 2015]
Mathis C, Bhattacharya T, Walker S (2015) The Emergence of Life as a First Order Phase Transition. Tech. rep., Arizona State University,
<http://arxiv.org/pdf/1503.02776>

- [McCarthy 1979]
McCarthy J (1979) Ascribing mental qualities to machines. In: Ringle M (ed) *Philosophical Perspectives in Artificial Intelligence*, Humanities Press, Atlantic Highlands, NJ, pp 161-195,
<http://www-formal.stanford.edu/jmc/ascribing/ascribing.html>
- [McCarthy and Hayes 1969]
McCarthy J, Hayes P (1969) Some philosophical problems from the standpoint of AI. In: Meltzer B, Michie D (eds) *Machine Intelligence 4*, Edinburgh University Press, Edinburgh, Scotland, pp 463-502,
<http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html>
- [Michie 1968]
D. Michie (1968), Memo Functions and Machine Learning, *Nature* vol 218, pp. 19-22, Nature Publishing Group,
<http://dx.doi.org/10.1038/218019a0>
- [Minsky 1987]
Minsky ML (1987) *The Society of Mind*. William Heinemann Ltd., London
- [Minsky 2006]
Minsky ML (2006) *The Emotion Machine*. Pantheon, New York (Text on his web page)
- [Penrose 1989]
Penrose R (1989), *The Emperor's New Mind: Concerning Computers Minds and the Laws of Physics*, OUP, Oxford
- [Penrose 1994]
Penrose R (1994), *Shadows of the mind: A Search for the Missing Science of Consciousness*. OUP, Oxford
- [Piaget, 1981]
Piaget, Jean, et al. 1981, *Possibility and Necessity Vol 1. The role of possibility in cognitive development*, U. of Minnesota Press, Tr. by Helga Feider from French in 1987,
- [Piaget, 1983]
Piaget, Jean, et al. 1983 *Possibility and Necessity Vol 2. The role of necessity in cognitive development*, U. of Minnesota Press, Tr. by Helga Feider from French in 1987,
- [Popper 1934]
Popper K (1934), *The logic of scientific discovery*. Routledge, London
- [Popper 1976]
Popper K (1976) *Unended Quest*. Fontana/Collins, Glasgow
- [Popper 1978]
Popper K (1978) Natural Selection and the Emergence of Mind. *Dialectica* 32(3-4):339-355,
<http://dx.doi.org/10.1111/j.1746-8361.1978.tb01321.x>
- [Powers 1973]
Powers WT (1973) *Behavior, the Control of Perception*. Aldine de Gruyter, New York
- [Rescorla 2015]
Michael Rescorla, Winter 2015, The Computational Theory of Mind, *The Stanford Encyclopedia of Philosophy*, Ed. Edward N. Zalta,
<http://plato.stanford.edu/archives/win2015/entries/computational-mind/>
- [Scott 2014]
Dana Scott, 2014, Geometry without points. (Video lecture, 23 June 2014, University of Edinburgh)
<https://www.youtube.com/watch?v=sDGnE8eja5o>

- [Schrödinger 1944]
Schrödinger E (1944). *What is life?* CUP, Cambridge
Extracts from this book, with some added comments, and a note added in the 1955 edition, can be found here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html>
- [Searle 1980]
J.R. Searle, 1980, Minds Brains and Programs, *The Behavioral and Brain Sciences*, 3, 3, (With commentaries and reply by Searle),
- [Seckbach & Rampelotto (2015)]
J. Seckbach and P. H. Rampelotto, 2015, Chapter 8: Polyextremophiles, in *Microbial Evolution under Extreme Conditions*, Ed. Corien Bakermans. DeGruyter pp 153-170
- [Senghas 2005]
Senghas A (2005) Language Emergence: Clues from a New Bedouin Sign Language. *Current Biology* 15(12):R463-R465,
<http://dx.doi.org/10.1016/j.cub.2005.06.018>
- [Shannon 1948]
Shannon C (1948) A mathematical theory of communication. *Bell System Technical Journal* 27:379-423 and 623-656
- [Simon 1967]
Simon HA (1967) Motivational and emotional controls of cognition. reprinted In: Simon HA (ed) in *Models of Thought*, Yale University Press, Newhaven, CT, pp 29-38
- [Sloman 1962]
Sloman A (1962) *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth* PhD thesis, Oxford University,
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962>
- [Sloman 1969]
Sloman A (1969) How to derive "better" from "is". *American Phil Quarterly* 6:43-52,
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1969-02>
- [Sloman 1970]
Sloman A (1970) "Ought" and "better". *Mind* LXXIX(315):385-394,
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1970-01>
- [Sloman 1971]
Sloman A (1971) Tarski, Frege and the Liar Paradox. In *Philosophy* 46(176):133-147,
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1971-03>
- [Sloman 1978]
Sloman A (1978) *The Computer Revolution in Philosophy. Philosophy, Science and Models of Mind*. Harvester Press (and Humanities Press), Hassocks, Sussex. Revised edition:
<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp>, 2015
- [Sloman(1992)]
A. Sloman, 1992. The Emperor's Real Mind, Review of Roger Penrose's *The Emperor's new Mind: Concerning Computers Minds and the Laws of Physics*, in *Artificial Intelligence*, 56, 1992, pp. 355--396,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#12>
- [Sloman 1983]
Sloman A (1983) Image interpretation: The way ahead? In: Braddick O, Sleight A (eds) *Physical and Biological Processing of Images* (Proceedings of an international symposium organised by The Rank Prize Funds, London, 1982.), Springer-Verlag, Berlin, pp 380-401,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#57>

[CogAff project]

The Birmingham Cognition and Affect Project: School of Computer Science, University of Birmingham, Birmingham, UK,

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

[Sloman 1993]

Sloman A (1993) The mind as a control system. In: Hookway C, Peterson D (eds) *Philosophy and the Cognitive Sciences*, Cambridge University Press, Cambridge, UK, pp 69-110,

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#18>

[Sloman 1995]

A. Sloman [1995], Exploring design space and niche space, in *Proceedings 5th Scandinavian Conference on AI*, Trondheim, 1995, IOS Press, Amsterdam,

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#41>

[Sloman 1996a]

Sloman A (1996a) Actual possibilities. In: Aiello L, Shapiro S (eds) *Principles of Knowledge Representation and Reasoning: Proc. 5th Int. Conf. (KR '96)*, Morgan Kaufmann Publishers, Boston, MA, pp 627-638,

<http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15>

[Sloman 1996b]

Sloman A (1996b) The SimAgent TOOLKIT - for Philosophers and Engineers (And Some Biologists, Psychologists and Social Scientists).

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

Sloman [IRREL]

Aaron Sloman, 2002, The irrelevance of Turing machines to AI, in *Computationalism: New Directions*, Ed. M. Scheutz, MIT Press, Cambridge, MA, pp. 87--127,

<http://www.cs.bham.ac.uk/research/cogaff/00-02.html#77>

[Sloman 2003]

Sloman A (2003) The Cognition and Affect Project: Architectures, Architecture-Schemas, And The New Science of Mind. Tech. rep., School of Computer Science, University of Birmingham, Birmingham, UK, (Revised August 2008)

<http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200307>

[Sloman 2006]

Sloman A (2006) Requirements for a Fully Deliberative Architecture (Or component of an architecture). Research Note, School of Computer Science, University of Birmingham, Birmingham, UK,

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/fully-deliberative.html>

[Sloman & Vernon, 2007]

Sloman, A and Vernon, D.(2007) A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots, 2007, Contribution to euCognition wiki,

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-requirements.html>

[Sloman 2008/2011]

Sloman A (2008/2011) Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)? Slide presentation, University of Birmingham. Installed 2008, updated 2011 and later.)

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>

[Sloman 2009]

Sloman A (2009) Architecture-Based Motivation vs Reward-Based Motivation. *Newsletter on Philosophy and Computers*, 09(1):10-13,

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html>

[Sloman 2010]

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

[Sloman 2011]

Sloman A (2011) What's information, for an organism or intelligent machine? How can a machine or organism mean? In: Dodig-Crnkovic G, Burgin M (eds) *Information and Computation*, World Scientific, New Jersey, pp 393-438,
<http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#905>

[Sloman 2011(b)]

Sloman A (2011) Evolved Cognition and Artificial Cognition
Some genetic/epigenetic trade-offs for organisms and robots
Or: Notes on how to do AI-Inspired Biology (AIIB)
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evolved-cognition.html>
Followup discussion paper after:
AISB workshop on AI-Inspired Biology (AIIB) Nottingham, 2010
<http://www.cs.bham.ac.uk/research/projects/cogaff/aiib/aiib.html>

[Sloman 2012]

Aaron Sloman (2012), Meta-morphogenesis and the Creativity of Evolution, in *Workshop on Computational Creativity, Concept Invention, and General Intelligence* Eds. T. R. Besold and K-U Kuehnberger and M. Schorlemmer and Alan Smaill, *Proc ECAI 2012* Montpellier, France.
<http://www.cs.bham.ac.uk/research/projects/cogaff/12.html#1203>

[Sloman 2013a]

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

[Sloman 2013b]

Aaron Sloman (2013b) Virtual machinery and evolution of mind (part 3) Meta-Morphogenesis: Evolution of information processing machinery.
Invited contribution to [Cooper and van Leeuwen 2013], pp 849-856,
<http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d>

[Sloman 2013c]

Aaron Sloman, (2013c) Meta-Morphogenesis and Toddler Theorems: Case Studies, Online discussion note, School of Computer Science, The University of Birmingham, <http://goo.gl/QgZU1g>

[Sloman 2015]

Aaron Sloman, (2015) Some (possibly) new considerations regarding impossible objects. Online discussion note, School of Computer Science, The University of Birmingham,
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

[Sloman 2017a]

Aaron Sloman (2017), Construction Kits for Biological Evolution, in *The Incomputable: Journeys Beyond the Turing Barrier*, Eds. S. Barry Cooper and Mariya I. Soskova, <http://www.springer.com/gb/book/9783319436678>

[Sloman and Chrisley 2003]

Aaron Sloman, Chrisley R (2003) Virtual machines and consciousness. *Journal of Consciousness Studies* 10(4-5):113-172,
<http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200302>

- [Sun 2006]
Sun R (2006) The CLARION cognitive architecture: Extending cognitive modeling to social simulation. In: Sun R (eds) in *Cognition and Multi-Agent Interaction*, Cambridge University Press, New York, pp 79-99,
<http://www.cogsci.rpi.edu/~rsun/sun.clarion2005.pdf>
- [Thompson 1917]
Thompson, D'Arcy W (1917) *On Growth and Form*. Cambridge University Press,
- [Trehub 1991]
Trehub A (1991) *The Cognitive Brain*. MIT Press, Cambridge, MA,
<http://people.umass.edu/trehub/>
- [Turing 1936]
Turing AM (1936) On computable numbers, with an application to the Entscheidungsproblem. *Proc London Math Soc* 42(2):230-265,
http://www.thocp.net/biographies/papers/turing_oncomputablenumbers_1936.pdf
- [Turing 1950]
Turing AM (1950) Computing machinery and intelligence. *Mind* 59:433-460, (reprinted in E.A. Feigenbaum and J. Feldman (eds) *Computers and Thought* McGraw-Hill, New York, 1963, 11-35)
- [Turing 1952]
Turing AM (1952) The Chemical Basis Of Morphogenesis. *Phil Trans R Soc London B* 237 237:37-72
Note: A presentation of Turing's main ideas for non-mathematicians can be found in Philip Ball, 2015, "Forging patterns and making waves from biology to geology: a commentary on Turing (1952) 'The chemical basis of morphogenesis'",
<http://dx.doi.org/10.1098/rstb.2014.0218>
- [Waddington 1957]
Waddington CH (1957) *The Strategy of the Genes*. MacMillan
- [Wagner 2014]
Wagner A (2014) *Arrival of the Fittest: Solving Evolution's Greatest Puzzle*. Published by: Oneworld Publications
- [Weir et al 2002]
Weir AAS, Chappell J, Kacelnik A (2002) Shaping of hooks in New Caledonian crows. *Science* 297(9 August 2002):981
- [Wiener 1961]
Wiener N (1961) *Cybernetics: or Control and Communication in the Animal and the Machine*. The MIT Press, Cambridge, MA, 2nd edition
- [Wilson 2015]
Wilson A (2015) *Metaphysical Causation*. Tech. rep., University of Birmingham, Department of Philosophy,
<http://alastairwilson.org/files/mc120715.pdf> video lecture:
<https://www.youtube.com/watch?v=j2l1wKrtlxs>
- [Wright 1997]
Wright I (1997) *Emotional agents*. PhD thesis, School of Computer Science, The University of Birmingham,
<http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#2>
- [Wright et al 1996]
Wright I, Aaron Sloman, Beaudoin L (1996) Towards a design-based analysis of emotional episodes. *Philosophy Psychiatry and Psychology* 3(2):101-126,
<http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#2>

File translated from Latex On 23 Jan 2016, 16:36, with a great deal of post-editing, which may have produced unnoticed errors.