
ESSENCE:
Evolution of Shared SEmaNtics in Computational Environments
A Marie Curie Initial Training Network

<https://www.essence-network.com/essence-events/summer-school/>
<https://www.essence-network.com/essence-events/summer-school/programme/>

Edinburgh University Informatics Department -- 24th-28th Aug 2015

NOTES FOR TUTORIAL
PRESENTED AT THE ESSENCE-2015 SUMMER SCHOOL
on
Evolved Construction Kits for Building Minds
(Including ideas about evolution of minds and mathematical abilities.)

The scientific/metaphysical explanatory role of construction kits:
fundamental and derived kits;
concrete, abstract and hybrid kits.

[Aaron Sloman](#)
[School of Computer Science, University of Birmingham.](#)

An extended abstract for the tutorial is available here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/essence-2015.html>


Video recording of the presentation (and other presentations)
available here:
<https://www.youtube.com/playlist?list=PLhz4qGm14yUDvzvLSRhmqBkotLyAYMzpd>

BACKGROUND:
The Turing-inspired Meta-Morphogenesis (M-M) project (2011--)
asks:
[How can a cloud of dust give birth to a planet](#)
[full of living things as diverse as life on Earth?](#)

In particular, what transitions in information processing
(contents of information, forms of information, mechanisms of
information, uses of information, sources of information, ... etc.
were required for production of all the varieties of life on Earth.

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>

A Protoplanetary Dust Cloud?



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

A partial answer is a theory of construction-kits, including construction kits produced by biological evolution and its products. This talk presents some preliminary, incomplete, ideas about types of construction-kit and their roles in biological evolution.

My aim is to inform and provoke potential future contributors to the project.

NOTE:

This is related to an invited paper for *The Incomputable* edited by S Barry Cooper and Mariya Soskova, to be published by Springer. That paper was frozen in Dec 2015, but related documents on this web site will continue growing, as the subject is so vast, and there are still so many gaps in our understanding (or mine, anyway).

NOTE (Added: 7 Jul 2016)

A video recording of the (somewhat disorganised) presentation is here (2hr45min):

<https://www.youtube.com/watch?v=RXL7cF-5VyU>

Including an attempt to explain the importance for biological evolution of parametric polymorphism.

NOTE:

This document was frozen in September 2015. It was originally derived from the main document on construction kits, which will continue to evolve here:

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

[JUMP TO TABLE OF CONTENTS](#)

This paper is available in two forms (HTML -- primary) and (PDF -- derived):

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

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

This tutorial presentation is a "frozen" version (produced 11 Sept 2015) of a web site that will continue to grow and change:

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

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

A closely related document can be found on slideshare.net (though that may become out of date):
<http://www.slideshare.net/asloman/construction-kits>

This document was prepared as a set of reconfigurable notes for a tutorial (presented at the ESSENCE Summer School, Edinburgh Informatics Dept, August 2015) for which the extended abstract is:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/essence-2015.html>

This is part of the Meta-Morphogenesis project, begun late 2011:

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

I did not realise until November 2014 (as a result of preparing a talk for a "Mind" Project workshop in Edinburgh), that a theory of evolved construction kits was required as an essential part of that project.

Background note in separate document, 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>

CONTENTS

ABSTRACT

0 Introduction: What is science? Beyond Popper and Lakatos

Now in a separate paper on

['Construction kits as explanations of possibilities'](#)

1 What biological possibilities need to be explained?

2 Fundamental and Derived Construction Kits (FCK, DCKs)

2.0 SMBC comic-strip comment on fundamentality

2.1 The variety of biological construction kits

2.2 More varied mathematical structures

3 Construction kits generate possibilities and impossibilities

3.1 Construction kits for making information-users

3.2 Different roles for information

3.3 Motivational mechanisms

3.4 Biological construction kits

4 Concrete (physical), abstract and hybrid (concrete+abstract) construction kits

-- Concrete kits:

-- Abstract kits:

-- Hybrid (abstract+concrete) kits:

4.1 Kits including external sensors and motors

4.2 Mechanisms for storing, transforming and using information

4.3 Mechanisms for controlling position, motion and timing

4.4 Combining construction kits

4.5 Combining abstract construction kits

5 Construction kits built during individual development

Note on relations with AI theories

(Added 13 May 2015)

Does ontogeny recapitulate phylogeny?

(First draft: Added 13 May 2015)

6 Some constructions exclude or necessitate others

6.1 Proof-like features of evolution

- [6.2 Euclid's construction kit](#)
- [6.2a Construction kits for internal languages](#)
- [6.3 Mathematical discoveries based on exploring construction kits](#)
- [6.4 Evolution's \(blind\) mathematical discoveries](#)
- [7 Varieties of Derived Construction Kit](#)
 - [7.1 A new type of research project](#)
 - [7.2 Construction-kits for biological information-processing](#)
 - [7.3 Representational blind spots of many scientists](#)
 - [7.4 Representing rewards, preferences, values](#)
- [8 Computational/Information-processing construction-kits](#)
 - [8.1 Infinite, or potentially infinite, generative power](#)
- [9 Types and levels of explanation of possibilities](#)

(This section needs to be clarified and reorganised.)

 - [9.1 Structural conformity](#)
 - [9.2 Process possibility](#)
 - [9.3 Process Abstraction](#)
 - [9.4 Grouping](#)
 - [9.5 Iterative or recursive abstraction](#)
 - [9.6 Self-assembly](#)
 - [9.7 Use of explicit intentions and plans](#)
 - [9.8 Construction kits linked across species](#)
 - [9.9 Concluding comment on explanations of possibilities](#)
- [10 Alan Turing's Construction kits](#)
 - [10.1 Beyond Turing machines: chemistry](#)
 - [10.2 Using properties of a construction-kit to explain possibilities](#)
 - [10.3 Bounded and unbounded construction kits](#)
[Quantum mechanisms and concurrency]

(Previously section 10.4: now in a separate document.)
- [11 Conclusion: Construction kits for Meta-Morphogenesis](#)

REFERENCES

FOOTNOTES

DOCUMENT HISTORY

ABSTRACT

Modified 24 Jan 2015; 31 Jan 2015; 19 Feb 2015; 18 Apr 2015

This is part of the Turing-inspired Meta-Morphogenesis project, introduced here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

The project aims to identify transitions in information-processing since the earliest proto-organisms, partly in order to provide new understanding of varieties of biological intelligence. This requires a study of related transitions: transitions in "construction-kits".

The initial "Fundamental Construction Kit" (FCK) is provided by the physical universe and the possibilities it supports, i.e. the physical and chemical structures and processes that it makes possible directly.

"Derived Construction Kits" (DCKs) are produced from the FCK by evolution, development, learning and culture.

Construction kits are of different types:

- *concrete*
- *abstract*
- *hybrid*.

Some are [meta-construction kits](#): able to create, modify or combine construction kits. Products of the construction kits are initially increasingly complex **physical** structures/mechanisms. Later products include increasingly complex **virtual** machines.

Philosophers and scientists have largely been ignorant about possibilities for virtual machinery until the development of computer systems engineering in the 20th Century introduced both new opportunities and new motivations for designing and building increasingly sophisticated types of virtual machinery, though 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 generally been noticed) may 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.

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.

This paper introduces a large research programme that seems to have a chance of being progressive, in the sense of [Imre Lakatos \(1980\)](#), rather than degenerative.

0 Introduction: What is science? Beyond Popper and Lakatos


Now in a separate paper on

['Construction kits as explanations of possibilities'](#)

1 What biological possibilities need to be explained?

We need to understand how the variety of mechanisms that existed at various stages in biological evolution combined with processes of natural selection proposed by Darwin, Wallace and others, made possible increasingly sophisticated organisms with enormous variation in size, physical form, environments, behaviours and later cognitive competences, as crudely depicted (with implied behaviours and types of information processing) here:

Evolutionary transitions from molecules to intelligent animals



These developments required not just changes in physical forms, but also changes in behaviour, and changes in information processing including cell division, epigenesis, control of physiological processes, control of actions, formation of cultures ... etc.

Within a biological species there can be considerable variation, not only across individuals, but even within an individual, over time. In the case of humans, there seem to be very rich forms of information-processing including implicit mathematical reasoning even before children can speak. See, for example, the "toddler theorems" illustrated in this document, including theorems about topological possibilities that appear to be intentionally used in pre-verbal behaviours:

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

Similar comments can be made about other intelligent species that never develop the abilities to discuss and prove the theorems they discover and use for example knot-tying weaver birds illustrated here:


<https://www.youtube.com/watch?v=6svAIgEnFvw>

Even before individual organisms develop mathematical capabilities, we'll see that natural selection implicitly acts as a "blind mathematician" discovering and using theorems about what is possible, for example in its production of homeostatic mechanisms (using negative feedback to maintain some state), as illustrated [below](#).

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>



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 drawing attention to that on her Facebook page.)

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). Some physicists have already proposed features of the FCK on this basis, including [Schrödinger, \[1944\]](#) and half a century later [Penrose, \[1994\]](#), among many others. (I think it will turn out that the connections are far less direct than they propose, because they have not investigated the many intermediate design requirements discussed or hinted at below.)

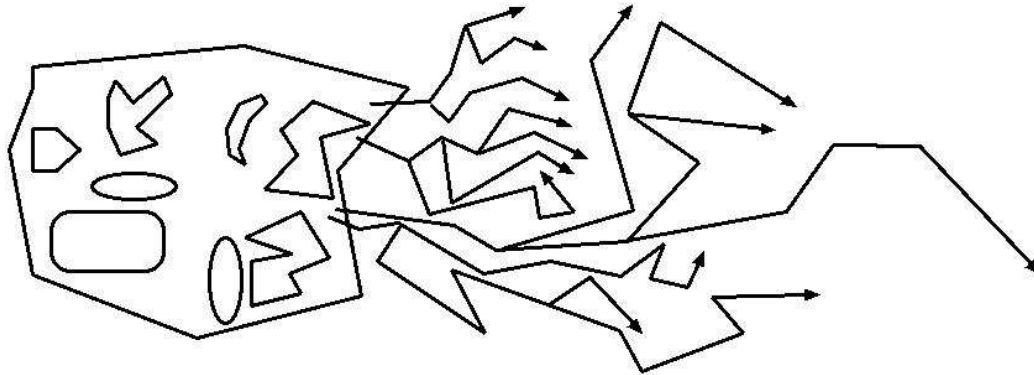
2 Fundamental and Derived Construction Kits (FCK, DCKs)

NB: everything written about construction kits here should be taken not as a developed theory, but a collection of loose ideas that serve to identify a (very difficult) long term research project, many fragments of which have begun to emerge from independently motivated research activities.

Life requires construction kits supporting construction of machines with many capabilities, including growing highly functional bodies, immune systems, digestive systems, repair mechanisms, and reproductive machinery. The requirements for life include information-processing (e.g. deciding what to repair) as well as physical construction (assembling matter).

The Fundamental Construction Kit (FCK)

The Fundamental Construction Kit (FCK) provided by the physical universe when our planet came into existence was sufficient to make possible all the forms of life that have so far evolved on earth, meeting challenges that drove selection of new life forms. The FCK also makes possible many unrealised but possible forms of life, in possible but unrealised types of physical environment. How does it make all these things possible?



The Fundamental Construction Kit (FCK) Evolutionary and other trajectories from the FCK through the space of possibilities

Figure FCK: Fundamental Construction Kit and possible trajectories

Think of time and increasing complexity going approximately from left to right.

Fig. [FCK](#), above indicates crudely how a common initial construction kit (FCK, on the left) could explain many possible trajectories in which components of the kit are assembled to produce new instances of possible living and non-living physical forms using increasingly complex mechanisms.

Products of a construction kit can have mathematical features that are useful, e.g. negative feedback. So evolution produces mathematical competences implicit in biological mechanisms. As explained (sketchily) below, this may lead later to explicit mathematical and meta-mathematical competences in some species, eventually providing new biological/evolutionary foundations for mathematics.

The history of technology, science and engineering includes many transitions in which new construction kits were derived from old ones. That includes the science and technology of digital computation, where new advances used (among other things):

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. many types of higher level programming language that are compiled to machine language or to intermediate level languages before programs start running;
10. 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.

Derived Construction Kits (DCKs)

Products of evolutionary trajectories from the FCK may combine to form Derived Construction Kits (DCKs) (some specified in genomes, and some designed or discovered, often unwittingly, by individuals, or groups), that speed up construction of more complex entities with new types of properties and behaviours, as crudely indicated in Fig. [DCK](#), below.




Figure DCK: Derived Construction Kit and new possible trajectories

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, new behaviours and new information-processing mechanisms. The shapes and colours (crudely) indicate qualitative differences between components of old and new construction kits. Time again goes (roughly) from left to right.

In cases of convergent evolution, new DCKs evolve in different species in different locations, with overlapping functionality, using different mechanisms. A DCK producing mechanisms enabling elephants to learn to use trunk, eyes, and brain to manipulate food may share features with a DCK enabling primates to acquire abilities to use hands, eyes, and brains to manipulate food. Both competences, apparently using related mathematical control structures, evolved after the last common ancestor.

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

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

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

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

Note on Making Possible:

The assertion "X makes Y possible" does not imply that if X does not exist then Y is impossible. All that is claimed is that one route to existence of Y is via existence of X. If X is built, that makes (deliberate or unplanned) construction of Y easier, faster, or in some cases more likely. However, other things than X can make Y possible, for instance, an alternative construction kit. So "makes possible" should be interpreted in our discussion as a relation of sufficiency, not necessity. The exception is the case where X is the FCK -- the *Fundamental Construction Kit* -- since all concrete constructions must start from that. If X and Y are abstract, it is not clear that there is something like the FCK to which they must be traceable. The space of abstract construction kits may not have a fixed "root" kit. However, the abstract construction kits that can be thought about by physically implemented thinkers may be more constrained.

Note on 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. So a physicist studying the FCK will need one or more (hybrid) construction kits for defining concepts, formulating questions, formulating theories and conjectures, constructing models, etc. Part of the process of science is extending the construction kit for theory formation, which includes extending the language used. Some of the later theories about DCKs (including theories about virtual machines in computer systems engineering) may include concepts that are not *definable* in terms of the concepts used in theories about the FCK, even though everything created using the DCK is *fully implemented* in the FCK. For more on this see [Sloman, \[2013a\]](#).

The concept of "Ontology" originally came from Metaphysics (e.g. Aristotle's work), but is now commonly used in science and engineering, especially information engineering. I think that the idea of "Descriptive Metaphysics" developed by Peter Strawson in [\(1959\)](#) is closely related, after some modification. This will be explained in more detail elsewhere.

2.1 The variety of biological construction kits

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

The simplest microbes use only a few (usually chemical) sensors providing information about internal states and the *immediate* external physical environment, and have very few behavioural options. They acquire, use and replace fragments of information, using the same types of internal information throughout their life.

More complex organisms acquire and use information about enduring spatial locations in extended terrain whose contents include static and changing resources and dangers, e.g. noxious substances or lurking predators. Some can construct and use complex (internal or external) information stores about their environment.

Some of them also acquire and use information about information-processing, in themselves and in others, e.g. conspecifics, predators and prey. What features of construction kits support these developments?

Some controlled systems have states represented by a fixed set of physical measures, often referred to as "variables" and "constants", representing states of sensors, output signals, and internal states of various sorts. Relationships between state-components are represented mathematically by equations, including differential equations, and possibly also constraints (e.g. inequalities) specifying restricted, possibly time-varying, ranges of values for the variables. Such a system with N variables has a state of a fixed dimension, N.

The only way to store new information in such "number-based" systems is in static or dynamic values for the variables -- changing "state vectors". A typical example is [Powers, \[1973\]](#), inspired by [Wiener, \[1961\]](#) and [Ashby, \[1952\]](#). There are many well understood special cases of this pattern, such as simple forms of homeostatic control using negative feedback. Neural net controllers may be very much more complex, with variables typically clustered into strongly interacting sub-groups, and perhaps groups of groups, etc. However mechanisms with this sort of mathematical structure are not the only ones used by natural selection.

2.2 More varied mathematical structures

In the last century, the variety of types of control in artefacts exploded, including use of logic, linguistics, and various parts of AI dealing with planners, learning systems, problem solving systems, vision systems, theorem provers, teaching systems, map-making explorers, automated circuit designers, program checkers, and many more.

The world wide web can be thought of as an extreme case of a control system made up of millions of constantly changing simpler control systems, interacting in parallel with each other and with millions of display devices, sensors, mechanical controllers, humans, and many other things. So the types of control mechanism in computer-based systems now extend far beyond the sorts familiar to control engineers, and studied in control theory.⁷

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

Numerical and non-numerical control

Many (though not all) human engineered control systems use numerical measures to represent states of whatever they are controlling, and the control mechanisms in such cases allow control interfaces to specify target numerical values for parts of the system or states and processes to be controlled.

However, many computational (information processing) control systems use non-numerical controls, e.g. specification of programs to be run, states to be achieved, tools to be used, etc. Likewise not all natural control functions are 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 the construction proceeds and the product (percept, thought, plan, map) becomes more complex -- unlike a fixed size collection of changing numerical values.

Different branches of numerical and non-numerical mathematics are suited to the problem of designing or understanding such systems, including graph theory, lattice theory, knot theory, category theory, set theory, logic, mathematical linguistics and others.

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

All this implies that traditional vector- and equation-based control theories, even with probabilistic extensions, are not general enough for intelligent control systems that build and use sentences, problem descriptions, changing ontologies, explanatory theories, plans of varying complexity, new types of learning mechanism, systems of motives, values, social rules, and rule-based games, among other things.

A fixed set of equations cannot adequately represent steady growth of increasingly complex molecular structures [[Anderson, 1972](#)]. Evolution, like human mathematicians and computer scientists millions of years later, built construction kits and information structures able to cope with structures and processes of changing complexity, unlike models and mechanisms based only on fixed sets of variables linked by equations -- unable to represent either the structure of the meaning of a complex sentence, such as this one, or what can exist on a skyscraper construction-site, or many other perceived processes, including waves breaking on a rocky seashore, an intricately choreographed ballet, or a symphony.

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

The huge variety of types of construction kit cannot be surveyed here. Instead of a complete theory: this paper merely presents a first-draft research framework within which gaps in our understanding can be discovered and in some cases filled, possibly over several decades, or even centuries. In particular, this first draft specifies some features of old and new construction kits, in the hope that additional research will extend the answers.

The planet on its own could not generate all those life forms. Energy from solar radiation is crucial for life on earth (though future technologies may remove that dependence). Other external influences that were important for the particular forms of life that evolved on earth included asteroid impacts, and cosmic radiation.⁸

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

According to some physical theories, every physical particle is (or can be) spread out over large areas, or possibly over the whole universe: nevertheless there must be differences in what exists in different places, for different processes can occur in different places. So the contents of the FCK are not necessarily distributed uniformly throughout the universe and some developments based on the FCK are impossible in certain parts of the universe lacking the required matter, or other pre-requisites.

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

As the earth formed, the new physical conditions created new DCKs that made the earliest life forms possible. [Ganti, \[2003\]](#) presents a deep analysis of requirements for a DCK that supports primitive life forms. That DCK (building on the FCK) made possible both the formation of pre-biotic chemical structures and very simple life forms, and also the *environments* in which they could survive and reproduce. But there's more to life than primitive life forms!

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

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

3 Construction kits generate possibilities and impossibilities

Explanations of how things are possible 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 clear boundaries.) Each kit makes possible certain types of construct, instances of which can be built by assembling parts provided in the kit. Some construction kits use products of products of biological evolution. For example, some birds' nests are assembled from twigs or leaves.

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

More generally, constructed objects can have features none of the components have, e.g. a hinge is a non-rigid object that can be made from rigid objects: two rigid objects with aligned holes through which a rod or screw is passed, creating a flexible object from non-flexible parts. A connected structure in a 2-D film cannot have a channel going right through it, whereas a 3-D structure can. There are many such examples of emergent novelty [[Anderson, 1972](#)]. I am not aware of any exhaustive taxonomy of ways of producing novel powers, structures and processes by combining old parts in new ways: apart from the implicit taxonomy in life forms.

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

supporting another that behaves as if hinged at the centre.

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

In some hybrid construction kits, such as chess, the physical pieces are not essential. For an expert, physical components of a chess set are dispensable: the abstract kit suffices for representing the abstract structures, states and processes, though communication of moves between players needs physical mechanisms, as does a player'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, which have different strengths and weaknesses.

3.1 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 construction kit using material that can be deformed under pressure, e.g. plasticine or damp clay. If some object, e.g. a coin, is pressed against a lump of the material the lump will change its shape, acquiring a new depressed portion whose surface has the inverted shape and size of part of the pressed object. Some entities with information-processing capabilities (e.g. archaeologists, or detectives) may be able to use the depression as a source of information about the coin. But the lump of material is not an information user. Likewise 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 portion of sand with footprints that last until the next time rain falls or a wind blows.

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

Additional mechanisms are required if available information is to be used: What sort of mechanism will depend on what sort 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). Additional mechanisms are required for recognising text, correcting spelling, analysing the structure of an image, interpreting it as a picture of a 3-D scene, or using information about the scene to guide a robot, or build a copy of the scene.

Different sorts of construction kit are required for producing those mechanisms. In organisms, the kits have different evolutionary histories: for example, mechanisms for finding, understanding, and correcting text evolved long after mechanisms able to use visual information for avoiding obstacles or for grasping objects.

In some cases, the mechanisms that use information seem to be direct products of biological evolution, including blinking as a defense mechanism, and other reflexes. In other cases, the detailed mechanisms are developed by individuals using mechanisms produced by evolution: for example: individual humans in different cultures develop different language-understanding mechanisms, but presumably they use a generic language construction kit shared with other humans. After use of such a kit begins it may be modified in ways that support further learning or development of a specific type of language. In [Chappell and Sloman, \[2007\]](#), the labels "preconfigured" and "meta-configured" were used for the contrast between direct specification of some feature in the genome and indirect specification, e.g. via use of an intermediate pre-specified mechanism for identifying problems and specifying solutions, or for providing parameters.

In some species, especially those using sexual reproduction, there may be considerable 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.

3.2 Different roles for information

Across all the diversity of biological construction-kits and the mechanisms that they produce in individuals there are some common recurring themes, including requirements for different types of information-based control state, such as information about how things actually are ("belief-like" information states), information about how things need to be for the individual information user ("desire-like" information states), and information about steps to take to achieve certain results ("procedural information states") -- See [Sloman, \[1996b\]](#). Biological construction kits can support those cases in different ways, depending on details of the environment, the animal's sensors, its needs, the local opportunities, and the individual's history. In some cases different mechanisms performing one of these functions share a common evolutionary precursor that has been modified in different ways. In other cases the mechanisms evolve independently -- convergent evolution.

A simple case is a thermostat that turns a heater on or off, discussed in [McCarthy, \[1979\]](#). It has two sorts of information: (a) about a target temperature set by a user (desire-like information) and (b) about current ambient temperature, provided by a sensor (belief-like information). The discrepancy between the two information items is used by the thermostat to select between turning a heater on, or off, or leaving it as it is. This is a very simple homeostatic mechanism, using information and a source of energy to maintain a state.

Many biological and human-designed control mechanisms acquire information through transducers and use the information in combination with energy sources, to produce, maintain or avoid various states of affairs. The causal role a physical state or change plays in controlling something else, e.g. controlling deployment of energy, altering direction of growth, selection of mode of analysis of information, among many others, can be described as providing information, in this case *control* information.

As [Gibson, \[1966\]](#) pointed out, acquisition of information often requires cooperation between processes of sensing and acting. In animal vision, saccades are actions that constantly select new information samples from the environment (e.g. from the optic cone). The use of that information is very different in different contexts, e.g. controlling grasping, controlling preparation for a jump, controlling avoidance actions, or sampling portions of text while reading. A particular sensor can therefore be shared between many control subsystems [[Sloman, 1993](#)], and the significance of particular sensor inputs will depend partly on which subsystems are in control of the sensor at the time, partly on which others happen to receive information from the sensor (assuming channels can be turned on or off).

The study of varieties of use of information in organisms is exploding, and now includes many mechanisms on molecular scales within much larger organisms as well as many intermediate levels of informed control, including sub-cellular levels (e.g. metabolism), physiological processes of breathing, temperature maintenance, digestion of food, blood circulation, control of locomotion, feeding and mating of large animals and coordination across communities, such as collaborative foraging in insects and trading systems of humans. Slime moulds include spectacular examples in which modes of acquisition and use of information change dramatically.⁹

The earliest evolved machines 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. In the simplest cases, local information is used immediately to select between alternative possible actions, as in a heating control, or trail-following mechanism. Uses of motion in haptic and tactile sensing and use of saccades, changing vergence, and other movements in visual perception all exemplify the interplay between sensing and doing, in "online intelligence".

But there are cases ignored by Gibson and by researchers opposed to cognitive theories, namely organisms that exhibit "offline intelligence", using perceptual information for tasks other than controlling immediate reactions, for example, reasoning about remote future possibilities or attempting to explain something observed. Offline intelligence requires use of 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.

One information-bearing structure (e.g. the impression of a foot, the shape of a rock or even a neural state) can provide very different information to different information-users, depending at least on (a) what kinds of sensors (including internal sensors) they can use to get information from the structure, (b) what sorts of information-processing (storing, analysing, comparing, combining, synthesizing, retrieving, deriving, using...) mechanisms the users have, (c) what sorts of needs or goals they can serve by using various sorts of information (knowingly or not).

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

For more on functions of vision and its connection with evolution of language see

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

Two Related Themes (intertwined): What are the functions of vision? How did human language evolve?

3.3 Motivational mechanisms

It is often assumed that every information user, U, must be trying to achieve some reward or avoid some punishment (negative reward). In that case, the effect of U acquiring some new item of information, I, will be to make some actions more likely, 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. Many AI systems and psychological theories are based on that assumption.

However, this ignores some of the sophistication of evolution. Animals are not all restricted to acting on motives selected on the basis of rewards expected by the individual. They may also have motive construction mechanisms that are simply triggered as "internal reflexes" by certain states of affairs, without having any knowledge or expectations regarding beneficial consequences of achieving those motives, 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 construction of new motives, which may or may not lead to behaviour, depending on how important other competing behaviours are. For example, in a kind person, watching someone fall may trigger a motive to rush to help. But that motive may not generate action if competing motives are too strong.

Moreover, such a motive need not be selected because acting on it will produce some reward for the actor, contrary to the widely held view that all motivation is reward-based. [Sloman, \[2009\]](#) labelled such reflex motive generation as "architecture-based motivation" in contrast with "reward-based motivation" where motives are selected on the basis of anticipated rewards. Behaviours apparently produced by architecture-based motivations can be observed in young children and the young of other playful intelligent animals. When watching such "idle" behaviours it may be tempting to invent hypothetical rewards but the assumption that expected rewards must always play a role in motive generation is just a prejudice. In some cases choosing between motives can take rewards into account, but moral principles or mere habits, may suffice instead.

One of the benefits of certain automatically triggered motives is that acting on them will sometimes produce new information, by sampling properties of the environment. That information may not be immediately usable, but in combination with other episodes of information storage may enable some later processes to analyse and reorganise the stored information. The individual need not have any conception of that later process when the information is acquired, though the ancestors of that individual may have benefited from the presence of the mechanisms of information gathering that were later used for information reorganisation (labelled "Representational Redescription" in [Karmiloff-Smith, \[1992\]](#)).

During evolution, and in some species 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).

If we don't understand the architecture and the many information-processing functions it supports, and how they are related, and how they grow and diversify, 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. The architectural knowledge about how such a systems works may not be expressible in sets of differential equations, or statistical learning mechanisms and relationships. (For important but partial attempts to characterise some architectural roles in human information-processing see [Minsky, \[1987,Minsky, \[2006,Laird et al. \[1987,Sun, \[2006\]](#). Compare [Sloman, \[2003\]](#). Earlier pioneering work was done by Herbert Simon, in "Motivational and emotional controls of cognition", reprinted in his *Models of Thought*, Yale University Press, pp. 29--38, 1967.)

The construction kits required for building information-processing architectures, with multiple sensors and multiple motor subsystems developing in complex and varied environments may differ in many ways, including:

- (a) what they provide as sources of information,
- (b) whether their mechanisms allow only immediate use of information or also allow storage for future use,
- (c) whether the information is used only in the form in which it is initially acquired or whether some is used after modification by mechanisms for analysing, parsing, interpreting, transforming, combining, or deriving information,
- (d) how long they can maintain information and whether it degrades with time,
- (e) what other types of information they can be combined with (possibly different kinds in different contexts),
- (f) whether use of the information requires additional information at various stages during the use (e.g. approximate information used to control a grasping action may require more precise information to be added in late stages of the grasp),
- (g) whether such additional information needs to have been acquired previously (like the combination of a lock, which is not needed during approach to the lock) or needs to be acquired from the environment while acting (like the precise locations of the lock's controls used in controlling hand movements),
- (h) whether outcomes of use of information can be used to modify previously acquired information (e.g. because the world has changed),
- (i) whether all the acquired information can be stored in the user, or whether external records are needed (e.g. diaries, filing systems, marks on trees),
- (j) whether the process of using acquired information can be terminated, temporarily suspended, or modified, by newly acquired, unexpected information,
- (k) whether information pathways through the system are fixed, or can be modified slowly by learning processes, or rapidly by context sensitive information management mechanisms,
- (l) whether only information known or expected to be true can be used or whether the organism can explore alternative hypothetical situations in order to work out their consequences (e.g. in constructing and using conditional plan steps),
- (m) whether successes and failures of processes using information merely cause adjustments in future actions or whether they can lead to re-assessment of the theories used (e.g. physical theories, chemical theories, theories about intentions of certain individuals, etc.) and in some cases major revisions of those theories,
- (n) whether surprising results can lead to modifications of the ontology used (e.g. adding new forces, new kinds of "stuff", genes, new quantum states, etc.).

An old example of ontology extension was the discovery of materials with magnetic properties. Others include electromagnetic fields, chemical elements and compounds, new life forms, quantum phenomena, and many more. There are probably many such ontology extensions during development of babies, toddlers and children, all of which remain un-noticed by adults.

This list of information-related differences between construction kits is not meant to be complete: it merely illustrates the complexity and variety of challenges in understanding the construction kits required for producing theories or models of biological information-processing. Not all biological information-processing systems have all these capabilities. Some capabilities are required for all organisms, though their forms can vary, for instance different mechanisms for maintaining a state by detecting and counteracting divergence from that state. So evolution of at least simple versions of those mechanisms must have happened very early.

Other information-processing functions, including abilities to acquire information about extended enduring environments external to the organism, and abilities to reason about hypothetical possibilities, and to modify ontologies used, are likely to have been products of relatively recent evolution, though I suspect little is known about which organisms have which capabilities, apart from a few species studied extensively by biologists.

Later developments derived from some of those abilities may have produced early types of mathematical knowledge and mathematical reasoning capabilities, such as reasoning about what is and is not topologically or geometrically possible in various situations in various situations.

There is much we still do not know about the construction kits used in these processes, and what they are used for. The Meta-Morphogenesis project aims to investigate the huge variety of uses of information and how they evolved, partly in the expectation that there will turn out to be many mechanisms and many information-uses that we have not noticed, that are essential for understanding, or replicating, the more complex control phenomena in living things, including brain functions.

I suspect that assumptions made by neuroscientists about the information-processing in brains omit some important types, and that AI researchers influenced by those assumptions therefore fail to replicate important functions of brains and minds in their machines, some discussed below.

Progress in this investigation may require major conceptual advances regarding what the problems are and what sorts of answers are relevant. E.g. "Where in the brain are discoveries made?" "Where do emotions occur in the brain?" "Where in the brain is musical ability?" "Where do visual experiences (qualia) occur in the brain?" "Where does understanding occur when you read a sentence?" are all nonsensical questions. (Compare "Where exactly in the USA was the president elected?") But that does not mean there are no mental states and processes, including detection of changes in qualia -- e.g. when internal self-monitoring processes notice changes in visual or other sensory information when nothing perceived in the environment has changed.

3.4 Biological construction kits

How did the FCK generate complex life forms? Is the Darwin-Wallace theory of natural selection the whole answer? Graham Bell wrote in [[Bell, 2008](#)]:

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

No: 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, in addition to evolved construction kits and their products.

The idea of evolution producing construction kits is not new, though they are often referred to as "toolkits". [Coates et al. \[2014\]](#) ask whether there is "a genetic toolkit for multicellularity" used by complex life-forms. Toolkits and construction kits normally have *users* (e.g. humans or other animals), whereas the construction kits we have been discussing (FCKs and DCKs) do not all need external users. [Ganti, \[2003\]](#) explained how chemistry supports self-sufficiency in very simple organisms.

Both generative mechanisms and selection mechanisms change during evolution (partly by influencing each other). Natural selection (blindly) uses the initial enabling mechanisms provided by physics and chemistry not only to produce new organisms, but also to produce new richer DCKs, including increasingly complex information-processing mechanisms. Since the mid 1900s, spectacular changes 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 whose products are

incomprehensible to thinkers familiar only with physics and chemistry, without any first hand experience of designing and implementing virtual machinery.

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, such as putting a shirt on a child¹⁰.

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 [7.2](#) below.)

4 Concrete (physical), abstract and hybrid (concrete+abstract) construction kits

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

Concrete Construction Kits (CCKs):

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

Abstract Construction Kits (ACKs):

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

The idea of a grammar is not restricted to verbal forms: it can be extended to many kinds of complex structures, e.g. grammars for sign languages, circuit diagrams, maps, architectural layouts and even molecules. Human sign languages use different structures from spoken languages.

A grammar does not specify a language: a *semantic* construction kit, structurally related to the grammar, is required for building possible *meanings* for the language to express. Use of a language depends on language users, for which more complex construction kits are required, also products of evolution and learning. (Evolution of various types of language is discussed in [Sloman, \[2008\]](#), which argues that internal languages must have evolved first, then sign languages.)

Hybrid (abstract+concrete) Construction Kits (HCKs):

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 is found.

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

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

4.1 Kits including external sensors and motors

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

Some construction kits allow assembly of such toys. More sophisticated kits include sensors that can be used to provide information for an internal mechanism that uses the information to take decisions concerning deployment of available energy, for instance using light, sonar, or in the case of rats, using whiskers, to gain information that allows frequent changes of direction or speed of motion, e.g. in order to avoid collisions, or in order to move towards a source of electrical or chemical energy when internal supplies are running low. Some examples are provided in [Braitenberg, \[1984\]](#), though he (or at least some of his admirers) unfortunately over-interpreted his vehicles as being capable of love, fear, etc.^{[11](#)}

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

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

4.2 Mechanisms for storing, transforming and using information

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

The offline/online use of visual information has caused much confusion among researchers, including attempts to interpret the difference in terms of "what" and "where" information.¹² Compare [Sloman, \[1983\]](#).

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

Architecture kits?

This raises the question whether evolution produced "architecture kits" able to combine evolved information-processing mechanisms in different ways, long before software engineers discovered the need. Such a kit could be particularly important for individuals that produce new subsystems, or modify old ones, during individual development, e.g. during different phases of learning by apes, elephants, and humans, as described in Section [5](#). (The BICA society aims to bring together researchers on biologically inspired cognitive architectures.¹³)

4.3 Mechanisms for controlling position, motion and timing

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

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

This contrasts starkly with the problems of getting software components available on a computer to interact: merely co-locating them in the same virtual machine on the same computer will not suffice. There are some rule-based systems 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 special "conflict resolution" mechanisms may be required if the conditions of more than one rule are satisfied.¹⁴

Spatial embedding of products allows new construction kits to be formed by combining two or more concrete kits. In some cases this will require modification of a kit, e.g. supporting combinations of lego and meccano by adding pieces with lego studs or holes alongside meccano sized screw holes. 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 the platform. In organisms, products of different construction kits may use complex mixtures of juxtaposition and adaptation. As mentioned in a separate paper, there is evidence that some organisms can also make use of non-local quantum effects when complex mechanisms are made of interacting components.

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

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

There are also mechanisms for controlling or varying rates of production of chemicals (e.g. hormones).

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

4.4 Combining construction kits

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

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

The application of functions to arguments is a very different process from assembling structures in space time. In the latter case inputs to the process form parts of the output, which need not be the case with a mathematical or computational function. If computers are connected via digital to analog interfaces, linking them to other things, e.g. surrounding matter, or if they are mounted on machines that allow them to move around in space and interact, that adds a kind of richness that goes beyond application of functions to arguments.

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

4.5 Combining abstract construction kits

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

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

Researchers on systems combining many kinds of functionality have found it useful to design information-processing *architectures* that provide frameworks for combining different mechanisms and information stores. This is particularly important in large projects where different research groups are working on sensors, learning mechanisms, motor subsystems, reasoning systems, motivational systems, various kinds of meta-cognition, etc., with associated sets of tools supporting processes of design, implementation, testing, debugging. Our own SimAgent toolkit [Sloman, \[1996c\]](#), mentioned in Note [14](#) is one among very many.

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

In addition to design patterns for physical mechanisms, biological evolution also discovered re-usable frameworks for assembling complex *information-processing* architectures, accommodating multiple interacting virtual machines, with different modifications developed by different species -- including humans [[Minsky, 1987](#),[Minsky, 2006](#)]. This is a topic for further research, which will provide new insights into complex mental states and processes, including forms of self-consciousness, varieties of affective states, and processes of cognitive development that help to explain mathematical development. ¹⁵

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

Creation of new construction kits may start by simply recording parts of successful assemblies, so that they can easily be reproduced. At later stages previous stores may be combined to form an appropriate "*meta-construction kit*" able to extend or modify or combine previously created construction kits. Evolution needs to be able to create new meta-construction kits using natural selection. Natural selection, the great creator/meta-creator, is now spectacularly aided and abetted by its products, especially humans and their products!

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

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

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

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

Multiple routes from genome to behaviours




Figure EVO-DEVO:

(Revised: 1 May 2015)

Figure derived from Chappell and Sloman, [2007].

A construction kit can give rise to very different individuals if the genome interacts with the environment in increasingly complex ways during development, allowing for enormously varied developmental trajectories based on the same genome. Precocial species use only the downward routes on the left, producing only "preconfigured" competences. Competences of altricial species, using staggered development, may be far more varied. Results of using earlier competences interact with the genome, producing "meta-configured" competences shown on the right.

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

Turing's (1952) examples of diffusing chemicals causing patterns when they interact include only formation of superficial 2-D patterns. Explaining the different ways in which features of a genome can directly or indirectly orchestrate many parallel processes of growth, development, formation of connections, etc. is a far greater challenge.

A possible framework for allowing abstract specifications in the genome to interact with details of the environment in instantiating complex designs is illustrated schematically in Fig. [EVO-DEVO](#). This generalises Waddington's "epigenetic landscape" metaphor [Waddington, \[1957\]](#), by allowing individual members of a species to partially construct their own epigenetic landscapes instead of merely following paths in a landscape that is common to the species. Related ideas are in [Karmiloff-Smith, \[1992\]](#).

Some of the implications of these ideas for attempts to understand genetic abnormalities such as autism are discussed in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/autism.html>

Note on relations with AI theories **(Added 13 May 2015)**

Many researchers in AI, psychology, philosophy and neuroscience have attempted to provide requirements and specifications of whole minds. My impression is that the vast majority focus on what they think is in an adult mind, and usually that means a "normal" adult human mind. Consequently they present their theories about the architecture of such a mind, without much consideration (if any) of either the processes of evolution by which such a mind could be the product of a succession of designs over millions of years with a variety of transitions, still mostly unknown, linking them.

There is a different sort of biological history that is generally ignored by those thinkers (at least the ones I have encountered, with a few exceptions such as Jean Piaget), namely the transitions in an individual mind between a fertilized egg (which most people I know would say has no mind) through a host of pre- and post-natal stages in which there are major qualitative transitions of many kinds. Although many developmental psychologists have attempted to probe some of the intermediate states and the transitions, they generally (through no fault of their own) lack the conceptual tools required for formulating an adequate theory of a *working* system: a theory that could in principle be used to specify the mechanisms in a machine that starts off either as an egg, or as some foetal stage, or as a new-born infant, and then constructs a succession of increasingly complex and varied mechanisms and information stores that in many cases go on developing throughout life. (This requires use of a [design-based approach](#) to theorising for which most educational systems fail to provide relevant competences. Being able to design a good theory is far more important than being able to test correlations for significance.)

During this process of development the body is also changing in the details of its shape, the sizes of parts, the strengths of the muscles, the types of posture and motion of which it is capable, and also in myriad internal ways, including development of neural mechanisms, of immune systems, of digestive mechanisms, or reproductive mechanisms, and many more.

Does ontogeny recapitulate phylogeny? **(First draft: Added 13 May 2015)**

Ernst Haeckel's idea that "ontogeny recapitulates phylogeny" (summarised in http://en.wikipedia.org/wiki/Recapitulation_theory) cannot be strictly correct. Yet at a high enough level of abstraction there may be some substance to the idea a developing organism goes through stages that loosely parallel the stages in evolutionary history of the species: for example both involve increasing complexity of physical structure and increasing differentiation of physical function. From our point of view there are also questions about the changes in information processing, including forms of control, forms of physical media used for storing information, types of use of information, types of information structure, types of information processing (e.g. use of information in control, use of information to derive new information, etc.), and types of internal communication of information within individuals.

Recapitulation is certainly too strong a claim, though there may be important similarities and analogies to be investigated. However, most AI theorists proposing information-processing mechanisms and architectures have focused on the adult form without much concern for earlier stages of development. An example is the work of a researcher who strongly influenced my own ideas when I was first

learning about AI, Marvin Minsky. A useful introduction to his ideas is provided by his former student Push Singh:

<http://web.media.mit.edu/~push/ExaminingSOM.html>

Examining the Society of Mind.

Minsky's later book *The emotion machine* built on those ideas. Both books are very rich stores of ideas about requirements for a human-like adult mind, and hints about the sorts of construction kits that may serve those requirements.

(To be continued)

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

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




Figure GAPS:

Interactions between structure and remaining possibilities:

If a rod that can swing about a point in a plane is in a gap, then the wider the gap the wider the possible swing, and the shorter the rod for a fixed size gap, the wider the possible swing. In general, interactions between structures and possibilities are more complex than this.

Moreover, what has been done may make some further steps possible and others impossible: e.g. the size of a gap between two rigidly assembled components will make it impossible to extend the structure by placing some components in the gap: A beam of 20cm square cross section cannot fit in a

10cm gap. Narrower beams can fit in the gap, but the angles by which their orientations can vary will depend on their diameter, the diameter of the gap, and other spatial relations. The narrower or shorter a beam in the gap, is the wider the angle through which it can rotate in a plane through the gap. The wider the gap is the wider the angle through which a beam of a certain width can rotate, while the longer the gap is the narrower the angle of rotation possible in that plane. Examples are in Figure [GAPS](#). Both human engineers and evolution can make use of similar, though usually more complex, mathematical relationships, in skeletal geometry for example.

Figure [Triangle](#) illustrates a different sort of example, where no physical properties of a structure (e.g. rigidity or impenetrability of materials) are involved, only spatial relationships. It presents a proof, found by Mary Pardoe, that internal angles of a triangle sum to a straight line, or 180 degrees.




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

The sequence of figures, demonstrates how the three-cornered shape has the consequence that summing the three angles necessarily produces half a rotation (180 degrees). Since the position, size, orientation, and precise shape of the triangle can be varied without affecting the possibility of constructing the sequence, this is a proof that generalises to any planar triangle. This is an unpublished proof reported to me by Mary Pardoe in the early 1970s.

Unlike the "standard" proofs, this proof makes no explicit reference to Euclid's parallel axiom. The human mathematical ability to look at a physical situation, or a diagram representing a class of physical situations, and reason about constraints on a class of possibilities sharing certain constraints may have evolved from earlier abilities to reason about changing affordances in the environment [[Gibson, 1979](#)]. Current AI perceptual and reasoning systems still lack most of these abilities, though that may change.

These are simple examples of the mathematical properties of construction kits (partly analogous to mathematical properties of formal deductive systems and AI problem solving systems).

As parts (or instances of parts) of the FCK are combined, structural relations between components of the kit have two opposed sorts of consequences: they make some further structures *possible*, and they make other structures *impossible* - and their absence or opposites, e.g. geometrical or topological properties, will then be *necessary* consequences of previous selection steps.

Note on modality

These examples illustrate how a construction kit with mathematical relationships can provide the basis for necessary truths and necessary falsehoods in some constructions, as in [Sloman, \[1962, Chap 7\]](#). See also Figure [Reutersvard](#) below. Such relationships between possibilities provide a deeper, more natural, basis for understanding modality (necessity, possibility, impossibility) than so called "possible world semantics". Being able to think about and reason about alterations in some limited portion of the environment is very common and a requirement for intelligent action [[Sloman, 1996a](#)]. In contrast

being able to think about the whole world, past, present and future, and the set of alternative complete worlds, is a far more demanding requirement. Moreover it is not clear how to decide whether an individual language user has that capability.

Since our examples of making things possible or impossible, or changing ranges of possibilities, are examples of causation, this also provides the basis for a Kantian notion of causation based on mathematical necessity [[Kant, 1781](#)], so that not all uses of the notion of "cause" are Humean (i.e. based on correlations), even if some are. Compare Section [6.3](#).¹⁶

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

6.1 Proof-like features of evolution

An unknown subset of the FCK, or perhaps a DCK or collection of DCKs, produced fortuitously as a side effect of formation of the earth, supported (a) primitive life forms and (b) processes of evolution that produced more and more complex forms of life, including new, more complex, derived, DCKs. New products of natural selection can make more complex products more reachable, as with toy construction kits, and mathematical proofs. Assembling a set of pre-built house parts (walls, door-frames, window-frames, etc.) provides routes to a collection of possible houses using those parts, where the routes are much shorter than routes starting from the primitive components. However starting from those parts will make some designs unreachable except by disassembling some of the parts first.

Moreover, there was not just one sequence of DCKs: 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.¹⁷

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

6.2 Euclid's construction kit

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

A different sort of geometry allows line segments to be translated and rotated in a plane while preserving their length. This is an assumption underlying the use of rulers for measuring length. Adding movable and rotatable line segments to Euclidean geometry allows an arbitrary angle to be divided into three equal parts, which is not possible in standard Euclidean geometry. See [Note 21\[a\]](#). A related construction is possible using "Origami geometry". The ability of humans to discover and explore such spaces of possibilities, may have played a role in the developments that led up to the discoveries assembled in Euclid's *Elements* ([Note 1](#)).

6.2a Construction kits for internal languages

These (proto-)mathematical abilities seem to have deep connections with more wide-spread animal abilities to detect and (implicitly?) reason about and make intelligent use of possibilities and impossibilities, abilities displayed, for instance, by squirrels, elephants, crows, and pre-verbal human toddlers, suggesting that the evolution of communicative uses of language by humans was preceded by more wide-spread evolution of powerful forms of representation and reasoning across a range of species [Sloman 2015a]. If that is correct, many arguments about the importance of uniquely human communicative languages in evolution of human intelligence may be at least partly mistaken.

Nevertheless it is true that there are very complex (and ill-understood) requirements for uses of (internal) languages for perceiving, reasoning, wondering whether, remembering, generalising, wanting, intending, planning, and controlling actions in accordance with plans or intentions. Significant subsets of those requirements that must have been met by evolutionary developments across a variety of non-human species before evolution of human communicative languages. In particular, the ability to perceive, want, intend, plan and execute complex intentions must have used highly structured internal forms of language with generative grammars for dealing with novelty and diversity in what is perceived, wanted, considered, intended, etc. These capabilities require the use of abstract construction kits such as grammars and mechanisms for creating and manipulating semantic contents, for internal languages, whose forms are not yet known, though they may be structurally more closely related to human sign languages than verbal languages, since sign languages obviously have more structures in common with actions than spoken languages do.

These conjectures may strike most readers as very strange and unfamiliar. The idea of spaces of possibilities generated by different sorts of physical construction kit (e.g. Lego or Meccano or Tinkertoy kits) may be easier for most people to understand than the comparison with generative powers of grammars or formal systems, though the two are closely connected, since grammars and axiom systems are both abstract construction kits that can be parts of hybrid construction kits.

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

Some challenges for construction kits used by evolution, and also challenges for artificial intelligence and philosophy, arise from the need to explain both how natural selection makes use of mathematical properties of construction kits related to geometry and topology, in producing organisms with spatial structures and spatial competences, and also how various subsets of those organisms developed specific topological and geometrical reasoning abilities used in controlling actions and solving problems, and finally how at least one species developed abilities to reflect on the nature of those competences and eventually, through unknown processes of individual development and social interaction, using unknown representational and reasoning mechanisms, managed to produce the rich, deep and highly organised body of knowledge published as Euclid's *Elements* (Note 1).

There are important aspects of those mathematical competences that as far as I know have not yet been replicated in Artificial Intelligence or Robotics¹⁸. I would argue that the results of statistical learning from previously acquired data that have recently produced impressive results in robots and AI software are seriously misleading because they will turn out to be dead ends when machines need the sorts of

(proto) mathematical creativity shown by many other species, or the problem-solving and creative designing capabilities of human engineers, architects, musicians, novelists, mathematicians, scientists, teachers and parents of adventurous children.

Why has it proved so difficult to replicate those competences? One reason may be that most of the detailed requirements have gone unnoticed, just as the rich mathematical structures of human languages, and animal visual competences, largely went unnoticed until the last century and a half. Is it possible that another problem is that currently understood forms of digital computation are inadequate for the tasks, whereas chemistry-based information-processing systems used in brains are much richer and more powerful, or even that there's some truth in speculations that quantum mechanisms play important roles in some aspects of animal intelligence? That question will be explored in another paper on requirements for construction kits used by natural selection. (Though I may lack the depth of understanding required for that task!)

Moreover, those who try to go too directly from hypothesized properties of the primordial construction kit to explaining advanced capabilities such as human self-awareness (e.g. [Schrödinger, \[1944\]](#), [Penrose, \[1994\]](#)) are likely to fail, because short-cuts will omit essential details of both the problems and the solutions, like mathematical proofs with gaps.

6.3 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 region of the surface, but that is impossible with only two lines, as noted in [Kant, \[1781\]](#). Likewise flat surfaces can be combined to enclose a volume, such as a tetrahedron or cube, but it is impossible for only three flat surfaces to enclose a finite space. It is not clear how humans detect such impossibilities: no amount of trying and failing can establish impossibility.

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

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

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

6.4 Evolution's (blind) mathematical discoveries

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

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

These observations seem to support a new kind of "Biological-evolutionary" foundation for mathematics (BEFM), 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 [5](#)).

As far as I know what is generally referred to as "foundations" by current mathematicians and philosophers of mathematics would not include BEFM. See the useful survey organised by Alexander Sakharov here: <http://sakharov.net/foundation.html>

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

These mathematical "discoveries" (discussed in more detail on the Meta-Morphogenesis web site [19](#)) started happening long before there were any humans doing mathematics (which refutes Wittgenstein's suggestion that mathematics is an anthropological phenomenon). Many of the discoveries were concerned with what is possible, either absolutely or under certain conditions, or for a particular sort of construction-kit.

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

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

For human evolution, a particularly important subclass of mathematical discoveries has been unwitting discovery and use of mathematical structures in the environment, a discovery process that starts in human children before they are aware of what they are doing, and in some species before uses of language for communication have developed. Examples are discussed in the "Toddler Theorems" document (Note [15](#)).

7 Varieties of Derived Construction Kit

Evolution and its products use the fundamental construction kit of physics and chemistry to produce *derived* construction kits, with new powers, including concrete, abstract and hybrid construction kits. DCKs may differ (a) at different evolutionary stages within a lineage, (b) across lineages (e.g. in different coexisting organisms such as plants, insects, vertebrates, etc.), and (c) during development of individuals that start from a single cell and develop mechanisms that support different kinds of growth, development and learning, providing new mechanisms for processing information, at different stages of development, discussed briefly in Section [5](#).

There is also variety in construction kits produced by cultures or ecosystems, illustrated by human languages, applied sciences as in bioengineering, notations for logic, the theory of computation and computer systems engineering. All 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 young child to acquire competence in use of any one of the thousands of different human languages (whether spoken or signed) in the first few years of life? There is evidence that children do not merely *learn* an existing language: 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](#)]. See also this video report <https://www.youtube.com/watch?v=pjtioIFuNf8>. There are also many possible human languages that might have developed but have not (yet).

Evolutionary trajectories leading to human spoken language capabilities may have gone from internal languages through collaborative actions then signed communication, then spoken communication, as argued in [[Sloman, \[2008\]](#)] and [[Sloman 2015a](#)].

If language acquisition were solely, or mainly, a matter of learning from language users, human languages could never have existed, since initially there were no expert users to learn from, and the process could not get started. This argument applies to many competences that might be thought to be based entirely on learning from experts, including mathematical expertise. So AI systems based on data-mining in samples of expert behaviours will never produce AI systems with human competences -- only subsets at best.

The history of computing since the earliest calculators demonstrates some of the kinds of change that can arise when new construction kits are developed. The technological changes were not merely changes of size, speed and memory capacity: there have been profound qualitative changes, in part because development of new layers of virtual machinery produced new types of mechanism, including new sorts of mutually interacting causal loops linking virtual machine control states with portions of external environments, for instance use of GPS in navigation using a "satnav" device.

Some of the new powers, states and processes include semantic contents referring to non-physical structures and processes, e.g. mathematical problems, rules of games, and mental contents including past or possible future mental contents and contents of other minds. So the machines cannot be *fully described* in the language of the FCK even though they are *fully implemented* in physical reality. (See [note on ontologies](#).)

We now understand some of the key components and modes of composition providing platforms on which *human-designed* layers of computation can be constructed, including subsystems closely but not rigidly coupled to the environment (e.g. using video cameras and propulsion by propellers, when coping with a cross-wind).

Several different sorts of "basic" abstract construction kits suffice to generate the forms of (discrete) computation so far studied. Those basic types include Turing machines, Post's production systems, Church's Lambda Calculus, and several more, each capable of implementing the others. 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.²⁰

One interpretation of the Church-Turing thesis is that these construction kits generate all possible forms of information-processing -- a claim I question. It is not obvious that those discrete mechanisms suffice for all biological forms of information-processing. In contrast, use of a wholly or partly chemical basis allows forms of computation that include both discrete and continuous mechanisms that were essential for some forms of biological assembly and information-processing. In some cases the assembly processes (including continuous changes such as folding, twisting, coming together, moving apart), seem to be self-controlling because partial structures constrain later possibilities. But the ability to form and release chemical bonds also provides discrete control. [Ganti, \[2003\]](#) shows how a chemical construction-kit supports forms of biological information-processing that don't depend only on external energy sources (a fact that's also true of battery-powered computers), and also supports growth and reproduction using internal mechanisms, which human-made computers cannot do (yet).

There may be many different sorts of construction-kit that allow different sorts of information-processing (computation) to be supported, including some that we don't yet understand. In particular, the physical/chemical mechanisms that support the construction of both physical structures and information-processing mechanisms in living organisms may have abilities not available in digital computers.²¹

7.1 A new type of research project

Most 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 explain the possibility of evolved biological DCKs. Most 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.

Among the physicists who have thought about this, Schroedinger ([1944](#)) pointed out that a theory of the physical basis of life should explain biological phenomena, 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 mentioned 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.²² An implication of this seems to be that information about the later form is present in the earlier forms: although no explicit use of information about how to make wings is apparent in the larval stage.

[Penrose, \[1994\]](#) attempted to show how features of quantum physics explain obscure features of human consciousness, especially mathematical consciousness, but ignored 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, must have built on previously evolved animal abilities, for instance abilities to see various types of affordance [\[Gibson 1979\]](#). The use of diagrams and spatial models by Penrose could be an example of that.

My impression is that when physicists attempt to explain features of human minds on the basis of their physical theory, they tend to try to jump too directly from aspects of fundamental physics to explanations of recently evolved, very complex human capabilities, such as mathematical capabilities in the case of Penrose.

But it is very unlikely that there are very abstract human mathematical abilities that somehow grow directly out of quantum mechanical aspects of the FCK, without depending on many intermediate developments, including the layers of perceptual, planning, and reasoning competences produced by billions of years of evolution. I have not yet fully understood Penrose's claims, however. Several other scientists have made related claims, including Stuart Hameroff, Henry Stapp, and many more. I'll later (tentatively) offer a different possible role for quantum mechanisms, that might explain aspects of the ability of human visual systems to take in and very rapidly organise, information about very complex new scenes, such as what is seen on turning a corner in a busy, unfamiliar city, or turning a corner in a richly stocked botanical center. Some of the ideas are under development in various papers on vision and a new separate paper here:

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

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.²³ However, the task can never be finished, since the process of construction of new derived construction kits may continue indefinitely, always producing more new kits with components and modes of composition that allow production of more complex types of *structure* and more complex forms of *behaviour* in organisms.

That idea is familiar to computer scientists and computer systems engineers since thousands of new sorts of computational construction kit (new programming languages, new operating systems, new virtual machines) have been developed from old ones in the last half century, making possible new kinds of computing system that could not previously be built from the original computing machinery, without introducing new intermediate layers, including, in some cases, 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 layers of virtual machinery in extending what information-processing systems can do, and the properties they can have. Evolution seems to have discovered that much earlier.

7.2 Construction-kits for biological information-processing

Applying the ideas from previous sections, we can speculate that the earliest evolved DCKs supported evolution of new physical/chemical mechanisms, soon followed by information-processing mechanisms used to gain benefits of selecting between available competences and tuning them -- on

the basis of results of perception, learning, motive formation, planning, and decision making. In some organisms, mathematical discovery processes, enabled production of competences used in generic understanding of sensory information, synthesis of separate information fragments into coherent wholes, and control systems using mechanisms for motive generation, plan construction, control of behaviour, and prediction.

Many of evolution's mathematical discoveries 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 use of new ontologies for separating situations requiring different behaviours.

Later still, construction kits used by evolution produced meta-cognitive mechanisms enabling individuals to notice and reflect on their own mathematical discoveries (enabling some of them to notice and remove flaws in their reasoning). In some cases those meta-cognitive capabilities allowed individuals to communicate their discoveries to others, discuss them, and organise them into complex highly structured bodies of shared knowledge, such as Euclid's *Elements* ([Note 1](#)). I don't think anyone knows how long all of this took, what the detailed evolutionary changes were, and how mechanisms of perception, motivation, intention formation, reasoning and planning 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. For a draft discussion of evolution of mathematical mechanisms, see [Note 24](#).

Many of these naturally occurring mathematical abilities have not yet been replicated in Artificial Intelligence systems or robots, unlike logical, arithmetical, and algebraic competences. Examples of topological reasoning about equivalence classes of closed curves not yet modelled in computers (as far as I know) are referenced in [Note 21](#). Even the ability to reason about alternative ways of putting a shirt on a child ([Note 10](#)) 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 a mixture! Perhaps there are important forms of representation or types of information-processing architecture still waiting to be discovered by AI researchers. 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. Compare [Penrose, \[1994\]](#).)

7.3 Representational blind spots of many scientists

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

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

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

Forms of mathematics normally 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 biology, and therefore for specifying what physicists need to explain.

Example: Many physicists, engineers and mathematicians who move into neuroscience assume that states and processes in brains need to be expressed as collections of numerical measures and their derivatives plus equations linking them, a form of representation that is well supported by widely used tools such as Matlab, but is not necessarily best suited for the majority of mental contents, and probably not even well suited for chemical processes where structures form and interact with multiple changing geometrical and topological relationships -- one of the reasons for the invention of symbolic chemical notations (now being extended in computer models of changing interacting molecular structures). Information-processing mechanisms also often need to manipulate non-numerical structures.

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

It is often assumed that all intelligent decision making uses positive or negative scalar reward or utility values that are comparable across options [[Luce and Raiffa, 1957](#)]. But careful attention to consumer magazines, political debates, and the varieties of indecision that face humans in real life shows that reality is far more complex. For example, many preferences are expressed in rules about how to choose between certain options. Furthermore preferences can be highly sensitive to changes in context. A crude example is the change in preference for type of car after having children. 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 use of comparisons of numerical values [[Sloman, 1969](#),[Sloman, 1970](#)]. Instead of a linear reward or utility metric, choices for intelligent individuals, or for natural selection, often involve a complex partial ordering network, with "annotated" links between nodes (e.g. "better" qualified by conditions: "better for", "better if" ...). In the Birmingham CogAff project [[Sloman, 2003](#)], those ideas later 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 and Sloman, 1993](#),[Beaudoin, 1994](#),[Wright et al. 1996](#),[Wright, 1977](#)]. Most AI/Cognitive Science models use much shallower notions of motivation.

Despite all the sophistication of modern psychology and neuroscience, I don't believe they currently have 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 human mechanisms that led to: mathematical discoveries over thousands of years, including mechanisms producing new 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 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, but that will not be pursued here.

8 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 in Section 2.

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.

8.1 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 (e.g. biological molecules with structures that withstand thermal buffeting, as required for genetic materials such as DNA) and also chemical locks and keys as in catalysis. Those mechanisms are essential for life as we know it, including forms of information-processing produced by evolution (mostly not yet charted). This topic is developed further in a separate paper (previously section 10.4 of this paper):

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

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, causation, that exist in our universe. However, physicists generally seek only to ensure that their construction kits are capable of accounting for phenomena observed in the physical sciences. Normally they do not assemble features of living matter, or processes of evolution, development, or learning, found in living organisms and try to ensure that their fundamental theories can account for those features also. There are notable exceptions mentioned above, such as Schrödinger and Penrose. Not all physicists who discuss physics and life (in my experience) attend to the many details of life, including the variety of forms it can take, the variety of environments coped with, the different ways in which individual organisms cope, the ways in which products of evolution become more complex and more diverse over time, and especially the many kinds of information-processing and control in individuals, in colonies (e.g. ant colonies), societies, and ecosystems.

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

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

If cosmologists and other theoretical physicists attempted to take note of 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 universe (or the Fundamental Construct Kit). In part that could be a consequence of the forms of mathematics known to physicists. (The challenge posed by [Anderson, \[1972\]](#) is also relevant: see Section [11](#), below.)

It may take many years of research to find out what exactly is missing from current physical theory that is required to explain biological phenomena. 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 examples of such challenges helps us assemble tests to be passed by future proposed theories: collections of possibilities that a deep physical theory needs to be able to explain.

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

Theories of utility measures may need to be replaced, or enhanced with new theories of how benefits, evaluations, comparisons and preferences, can be expressed. We must also avoid assuming optimality. 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, because they are lucky, at least for a while. Likewise human decision making.

9 Types and levels of explanation of possibilities **(This section needs to be clarified and reorganised.)**

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 such cranes 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. The last feature may be relevant to modes of specification of constructions either in a genome or in a learnt or invented specification for a solution to a type of problem.

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: merely as a first experimental foray, preparing the ground for future work:

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

Some structures are specified in terms of piece-wise relations in a language with grammatical structures and compositional semantics. However in such languages it is often possible to specify parts and relations of a complex structure 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". A similar phenomenon can occur in non-verbal forms of representation. For example, complex structures made of perfectly possible fragments with perfectly possible piece-wise relations may be impossible as wholes, and many such impossible 3-D structures have been depicted in drawings and paintings, e.g. in pictures of impossible objects and scenes by Hogarth, Reutersvard, Escher, Penrose, and others. See Figure [Reutersvard](#) below.




Figure REUTERSVARD:

Pictures depicting possible and impossible scenes.
 (Picture on right by Oscar Reutersvard 1934)

Some powerful representational construction kits can depict things that cannot possibly exist, like the configuration on the right. See also

http://en.wikipedia.org/wiki/Impossible_object

Essentially this idea was already known to William Hogarth, who produced an engraving in 1754, entitled "Satire on False Perspective", analysed in

http://en.wikipedia.org/wiki/Satire_on_False_Perspective. A fairly good quality version is here (see 'Other versions' section):

<http://en.wikipedia.org/wiki/File:Hogarth-satire-on-false-perspective-1753.jpg> . For more examples and references see http://en.wikipedia.org/wiki/Impossible_object.

Similar examples can occur in arithmetical descriptions. Consider this specification of a number:

A number greater than 23 and less than 29 which has no divisors other than 1 and itself and is greater than any other number that is between 24 and 28 inclusive.

That describes a number which, among other things, is between 23 and 29, is distinct from those two, and is prime. But there is not and cannot be such a prime number, since all the eligible numbers, namely 24,25,26,27 and 28, have proper factors. So pictures of impossible objects are a special case of a more general phenomenon.

These examples show that in a complex specification of some entity *structural conformity* can be local or global. Local structural conformity with rules of grammar or geometric constraints can hold in many overlapping subsets of a description, picture or specification, even though the whole thing thus represented is impossible. (It is arguable that the contents of many religious beliefs are like this.)

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 [[Sloman, 1971](#)].

Of course, if some complex product of a construction kit is not merely specified but actually constructed, then that demonstrates conclusively that it is possible, unlike the construction of a description or depiction of the product. (Something like this idea underlies "constructivist" philosophy of mathematics. Compare the next item.)

9.2 Process possibility: The second type of answer to "How is X possible?" demonstrates constructability by describing a sequence of spatial trajectories by which the required collection of parts could be assembled. This may include processes of assembly of temporary supports to hold parts in place before the connections have been made that make them self-supporting or before the final supporting structures have been built (as often happens in large engineering projects, such as bridge construction).

In some cases, 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. Different routes to the same end result may differ only trivially (e.g. in the order in which two unrelated changes are made) or in more complex ways, e.g. requiring different temporary supports.

In some cases a complex object has type 9.1 possibility (structural conformity) but not type 9.2 (process possibility). For example, from a construction kit containing several rings it is possible to assemble a *pile* of three solid, rigid, impermeable rings, but not possible to assemble a *chain* composed of those rings even though each of the parts of the chain is exactly like the parts of the pile. The chain of linked rings can be described, even though it cannot be assembled: no possible construction process can be described, since parts of the rings cannot move through parts of other rings. Of course, the construction may be possible using a kit whose components are simpler than complete rings, and from which such rings can be made. E.g. linked rings could be assembled using plasticine, or a suitable chemical construction kit.

9.3 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 a number of 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 *optimal* solutions. So they design mechanisms that search using an optimisation process, ignoring the possibility of mechanisms that can find sets of possible solutions (e.g. routes) initially considered as a class of *equivalent* options, leaving questions about optimal assembly to be settled later, if needed. These remarks are closely related to the origins

of abilities to reason about geometry and topology.²⁵

9.4 Grouping: Another form of abstraction is related to the difference between **9.1** and **9.2**. If there is a sub-sequence of assembly processes, whose order makes no difference to the end result, they can be grouped to form an unordered "composite" move, containing an unordered set of moves. If N components are moved from initial to final states in a sequence of N moves, and it makes no difference in what order they are moved, merely specifying the set of N possibilities without regard for order collapses N factorial sets of possible sequences into one composite move. If N is 15, that will collapse 1307674368000 different sequences into one. If each move can be represented only by start and end states, as in **9.3**, that will further reduce the space of alternatives.

Sometimes a subset of moves can be made in parallel. E.g. someone with two hands can move two or more objects at a time, in transferring a collection of items from one place to another. Parallelism is particularly important in many biological processes where different processes occurring in parallel constrain one another so as to ensure that instead of all the possible states that could occur by moving or assembling components separately, only those end states occur that are consistent with parallel constructions. 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, if completing some too soon may make the required final configuration unachievable. For example, rings being completed before being joined could prevent formation of a chain.

9.5 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.

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 Section [5](#) on epigenesis.)

9.6 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).

9.7 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, 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.

9.8 Construction kits linked across species: Some construction kits involve cross species relationships, including predator-prey relationships, parasite-host relationships and symbiotic relationships (including domestication of one species by another). In all these cases there are (at least?) two types of organism each using an evolved construction kit where the kits evolved either as a result of benefits of cooperation between the types or as a result of competition between the types (so-called evolutionary "arms races").

http://en.wikipedia.org/wiki/Evolutionary_arms_race

There are similar relationships between construction kits used by males and females, including collaborative construction kits supporting cooperation to achieve fertilization and in some cases competitive construction kits supporting competition between males and females, insofar as some features allow males to increase the number of their offspring by increasing the burden on females (e.g. evolved behaviour patterns requiring females to do most or all of the rearing as well as bearing of young) and others work in the opposite direction by requiring males to do most of the food-gathering, or even most of the care for offspring. An extreme case occurs in the seahorse: the male carries fertilized eggs in a pouch until they are ready to be released to fend for themselves.

<http://animals.nationalgeographic.com/animals/fish/sea-horse/>

9.9 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 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 young. 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 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 changes in types of computer programming language and software design specification languages. Evolution may have beaten us to important ideas, by millions of years!

That all these levels of abstraction are possible is a metaphysical feature of the universe, implied by the generality of the FCK.

10 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 in which symbols can be inserted.

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

Another type of construction kit with related properties is Conway's Game of Life²⁶, a construction kit that creates changing patterns in 2D regular arrays. Stephen Wolfram has written a great deal about the diversity of constructions that can be explored using such cellular automata. Neither a Turing machine nor a Conway game has any external sensors: once started they run according to their stored rules and the current (changing) state of the tape or grid-cells. In principle either of them could be attached to external sensors that could produce changes to the tape of a Turing machine or the states of some of the cells in the Life array. However any such extension would significantly alter the powers of the machine, and theorems about what such a machine could or could not do 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 (from keyboards, buttons, joy-sticks, mice, electronic piano keyboards, or network connections) and can also send signals to external devices. Theorems about disconnected Turing machines may not apply to machines with rich two-way interfaces to an 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. 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.

(It is perhaps worth mentioning that one of the differences between 2-D and 3-D structures is that a connected 3-D structure can have an interior space and an exterior space with two or more distinct routes joining them (essential for an organism to consume food through one opening and excrete through a separate one), whereas in a 2-D space any structure with two holes (or a through-route) would not be fully connected. This severely limits the possibilities for 2-D life forms.)

10.1 Beyond Turing machines: chemistry

Turing also explored other sorts of construction kits, including types of neural nets and extended versions of Turing machines with "oracles" added. 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 I am sure many who have played with construction kits will have noticed). On the other hand if the kit allows arbitrarily small changes it will, in principle, permit exhaustive searches in some sub-spaces. The exhaustiveness comes at the cost of a very much larger (infinite, or potentially infinite!) search-space. That feature could be useless, unless the space of requirements has a structure

that allows approximate solutions to be useful. In that case a mixture of big jumps to get close to a good solution, followed by small jumps to home in on a (locally) optimal solution can be very fruitful: a technique that has been used by Artificial Intelligence researchers, called "simulated annealing".²⁷

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 to be found with *relatively* little searching (compared with other search spaces), after which small random changes allow improvements to be found. I have not yet read the book but it seems to illustrate the importance for evolution of the types of construction-kit available.²⁸ I have not yet had time to check 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.

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

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

10.3 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, or whether certain pairs of numbers can be adjacent, whether the same number can occur in more than one square. For a given grid and a given set of constraints 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 kits 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.

10.4 More on Quantum Mechanisms

Now in a separate document

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

11 Conclusion: Construction kits for Meta-Morphogenesis

As I was finishing a first draft of this paper I found a useful survey by Evelyn Fox Keller (briefly summarised [here](#)), of previous attempts to show how life and its products relate to the physical world: [Keller, \[2008\]](#), [Keller, \[2009\]](#). She 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*". This paper outlines a project to do that and more: including branching layers of new *derived* construction kits produced by evolution, development and other processes. The physical world clearly provides a very powerful (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, human designed, construction kits 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. The idea is still at an early stage of development. 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.

In particular, construction-kits for building virtual machinery can help to explain how minds and their contents can exist in a material world, with causal powers that affect the material world. 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 et al, \[in press\]](#) a distinction is made between two "principled" options for the relationship between the basic constituents of the world and their consequences. In the "Humean" option there is nothing but the distribution of structures and processes over space and time, though there may be some empirically discernible patterns in that distribution. The second option is "modal realism", or "dispositionalism", according to which there is something about the primitive stuff and its role in space-time that constrains what can and cannot exist, and what types of process can or cannot occur. This paper supports a "multi-layer" version of the modal realist option (developing ideas in [Sloman, \[1962, Sloman, \[1996a, 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 [\[Anderson, 1972\]](#), namely how we should understand the relationships between different levels of complexity in the universe (or in scientific theories). The reductionist alternative claims that when the physics of elementary particles (or some other fundamental physical level) has been fully understood, everything else in the universe can be explained in terms of mathematically derivable consequences of the basic physics. Anderson contrasts this with the anti-reductionist view that different levels of complexity in the universe require "entirely new laws, concepts and generalisations" so that, for example, biology is not applied chemistry and

psychology is not applied biology. He writes: "Surely there are more levels of organization between human ethology and DNA than there are between DNA and quantum electrodynamics, and each level can require a whole new conceptual structure". However, the structural levels are not merely in the concepts used by scientists, but actually in the world.

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

In future, physicists wishing to show the superiority of their theories, should attempt to demonstrate mathematically and experimentally that they can explain more of the potential of the FCK to support varieties of construction kit required for, and produced by, biological evolution than rival theories can. Will that be cheaper than building bigger better colliders? Will it be harder?

Construction kits are generative: They explain possibilities

A construction kit explains the possibility of a set of possible construction processes, with mathematical properties and limitations. Evolution and development demonstrate new possibilities for construction kits: evolution as a "blind theorem prover", proving "theorems" about what is and is not possible for the kits used.

The requirement to show how the FCK makes *everything else* possible provides a challenge for physicists: demonstrate that the fundamental theory can explain how all the products of natural selection are possible. A core thread is the connection of control and semantic information. The aim is to explain, not reduce.

Endnote

In 1946 Turing wrote to W. Ross Ashby 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):

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

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

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.

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 W.R. (1952) Chapman and Hall, London

[\[Beaudoin 1994\]](#)

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

[\[Braitenberg 1984\]](#)

Braitenberg V (1984) *Vehicles: Experiments in Synthetic Psychology*. The MIT Press, Cambridge, MA

[\[Chappell and Sloman 2007\]](#)

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>

[\[Coates et al 2014\]](#)

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

[\[Esfeld et al in press\]](#)

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>

[\[Ganti 2003\]](#)

Ganti T (2003) *The Principles of Life*. OUP, New York, Eds. Eörs Szathmáry & James Griesemer, Translation of the 1971 Hungarian edition

[\[Gibson 1966\]](#)

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

[\[Gibson 1979\]](#)

Gibson JJ (1979) *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA

[\[Hameroff and Penrose 2014\]](#)

Hameroff S, Penrose R (2014) Consciousness in the universe: A review of the 'orch or' theory. *Physics of Life Reviews* 11(1):39--78, DOI <http://dx.doi.org/10.1016/j.plrev.2013.08.002>

[\[Jablonka and Lamb 2005\]](#)

Jablonka E, Lamb MJ (2005) *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. MIT Press, Cambridge MA

[\[Kant 1781\]](#)

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

[\[Kauffman 1995\]](#)

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

[\[Keller 2008\]](#)

Keller, Evelyn Fox (2008) Organisms, Machines, and Thunderstorms: A History of Self-Organization, Part One. *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, Evelyn Fox (2009) Organisms, Machines, and Thunderstorms: A History of Self-Organization, Part Two. Complexity, Emergence, and Stable Attractors. *Historical Studies in the Natural Sciences* 39(1 (Winter)):1-31, <http://www.jstor.org/stable/10.1525/hsns.2009.39.1.1>

[\[Laird et al 1987\]](#)

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 RD, Raiffa H (1957) John Wiley and Sons, Inc.; Chapman and Hall, New York; London

[\[McCarthy 1979\]](#)

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>

[\[Minsky 1987\]](#)

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

[\[Minsky 2006\]](#)

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

[\[Penrose 1994\]](#)

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

[\[Popper 1934\]](#)

Popper K (1934) The logic of scientific discovery. Routledge, London

[\[Powers 1973\]](#)

Powers WT (1973) Behavior, the Control of Perception. Aldine de Gruyter, New York

[\[Schrödinger 1944\]](#)

Schrödinger E (1944) What is life? CUP, Cambridge

A commented subset of the book is online here:

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

[\[Senghas 2005\]](#)

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>

[\[Sloman 1962\]](#)

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/07.html#706>

[\[Sloman 1969\]](#)

Sloman A (1969) How to derive "better" from "is". American Phil Quarterly 6:43-52, <http://www.cs.bham.ac.uk/research/cogaff/papers.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/05.html#200508>

[\[Sloman 1971\]](#)

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

[\[Sloman 1978\]](#)

Sloman A (1978) The Computer Revolution in Philosophy. Harvester Press (and Humanities Press), Hassocks, Sussex, <http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp>

[\[Sloman 1983\]](#)

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/06.html#0604>

[\[Sloman 1993\]](#)

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>

[\[Sloman 1996a\]](#)

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) Beyond turing equivalence. In: Millican P, Clark A (eds) *Machines and Thought: The Legacy of Alan Turing (vol I)*, The Clarendon Press, Oxford, pp 179-219, <http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#1>, (Presented at Turing90 Colloquium, Sussex University, April 1990)

[\[Sloman 1996c\]](#)

Sloman A (1996c) 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>

[\[Sloman 2000\]](#)

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

[\[Sloman 2003\]](#)

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, <http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200307> (Revised August 2008)

[\[Sloman 2006\]](#)

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

[\[Sloman 2008\]](#)

Sloman A (2008) Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)? <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang>

[\[Sloman 2009\]](#)

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>

[\[Sloman 2013a\]](#)

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

[\[Sloman 2013b\]](#)

Sloman A (2013b) Virtual machinery and evolution of mind (part 3) meta-morphogenesis: Evolution of information-processing machinery. In: Cooper SB, van Leeuwen J (eds) *Alan Turing -- His Work and Impact*, Elsevier, Amsterdam, pp 849-856, <http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d>

[\[Sloman 2015a\]](#)

Sloman A (2015a), Two Related Themes (intertwined): What are the functions of vision? How did human language evolve? (Online Presentation, School of Computer Science, University of Birmingham) <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>

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

[\[Sun 2006\]](#)

Sun R (2006) The CLARION cognitive architecture: Extending cognitive modeling to social simulation. In: Sun R (ed) *Cognition and Multi-Agent Interaction*, Cambridge University Press, New York, pp 79-99, <http://www.cogsci.rpi.edu/~rsun/sun.clarion2005.pdf>

[\[Turing 1936\]](#)

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

[\[Waddington 1957\]](#)

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

[\[Wiener 1961\]](#)

Wiener N (1961) *Cybernetics: or Control and Communication in the Animal and the Machine*. The MIT Press, Cambridge, MA, 2nd edition

[\[Wright 1977\]](#)

Wright I (1977) *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, Sloman A, 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#22>

FOOTNOTES:

¹ Euclid's *Elements*

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

²

<http://plato.stanford.edu/entries/democritus/#2>

<http://en.wikipedia.org/wiki/Democritus>

³ Chapter 2:

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

⁴ Extended in

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

⁵

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

Explaining possibilities.

⁶ <https://www.youtube.com/watch?v=wcXSpXyZVuY>

⁷ See http://en.wikipedia.org/wiki/Control_theory

http://en.wikipedia.org/wiki/Nonlinear_control

⁸ The role of entropy is discussed briefly in:

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

⁹

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

¹⁰ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/shirt.html>

¹¹ <http://www.it.bton.ac.uk/Research/CIG/Believable%20Agents/>

¹² http://en.wikipedia.org/wiki/Two-streams_hypothesis

¹³ Some examples are here:

<http://bicasociety.org/cogarch/>

¹⁴ The Birmingham SimAgent toolkit is an example
<http://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html>

¹⁵ 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. Deep insights come from individual developmental trajectories rather than statistical patterns of development across individuals.)

¹⁶ 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>

¹⁷ <http://en.wikipedia.org/wiki/Symbiogenesis>

¹⁸ Some of them listed in
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/mathstuff.html>

¹⁹
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html#blind-theorem>

²⁰ For more on this see:
http://en.wikipedia.org/wiki/Church-Turing_thesis

²¹ Examples of human mathematical reasoning in geometry and topology that have, until now, resisted replication on computers are presented in these discussion papers (and others referenced therein):
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html>

^{21[a]} Compare this discussion of angle trisection with the examples in [Note \[21\]](#) above:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html>

²² <http://en.wikipedia.org/wiki/Pupa>
<http://en.wikipedia.org/wiki/Holometabolism>

²³ <http://en.wikipedia.org/wiki/J.\ B.\ S.\ Haldane>

²⁴
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/befm-sloman.pdf>

²⁵ 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>

²⁶ <http://en.wikipedia.org/wiki/Conway.27s.Game.of.Life>

²⁷ One of many online explanations is
<http://www.theprojectspot.com/tutorial-post/simulated-annealing-algorithm-for-beginners/6>

²⁸ An interview with the author (Wagner) is online at
<https://www.youtube.com/watch?v=wyOgCMZdv6E>

File translated from T_EX by T_TH, version 4.05.
On 16 Apr 2015, 23:44, then subsequently edited by hand, by A.Sloman

DOCUMENT HISTORY

Begun: 15 Dec 2014

(Based partly on earlier documents on [the Meta-Morphogenesis project](#))

Last updated:

11 Sep 2015 Updated version on www.slideshare.net:

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

27 Aug 2015: Version of [construction-kits.html](#) saved for this ESSENCE tutorial.

21 May 2015: Section 10.4 on quantum mechanics now moved to a separate document:

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

18 May 2015: tidied up and expanded, including discussion of quantum mechanisms.

13 May 2015: started adding a section on "[relations with AI theories](#)" and whether "[ontogeny recapitulates phylogeny](#)" in information processing mechanisms.