



(DRAFT: Liable to change)
Last updated: 7 Apr 2022

Biologically Evolved Forms of Compositionality **Structural relations and constraints** **vs Statistical correlations and probabilities**

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

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

Short link here: <https://goo.gl/hQikAJ>

Aaron Sloman

School of Computer Science, University of Birmingham, UK

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

Some of the key ideas included in this work
were developed in collaboration with Jackie Chappell:

<https://www.birmingham.ac.uk/staff/profiles/biosciences/chappell-jackie.aspx>

Note: 7 Apr 2022

Recent additions regarding evolution of compositionality extend these ideas.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-CCB-2022.html>

Notes related to a talk given on 20 Sep 2018,
at the First Symposium on Compositional Structures (SYCO 1)
School of Computer Science, University of Birmingham, UK

<http://events.cs.bham.ac.uk/syco/1/>

This is part of [The Meta-Morphogenesis project](#)

My talk at SYCO-1 used a much shorter version of this paper namely:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/compositionality-syco-1.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/compositionality-syco-1.pdf>

Also available as: <http://goo.gl/DjDQNS>

CONTENTS

[Keywords/Phrases](#)

[Preface: Evolution uses increasingly abstract evolved and developed types](#)

[Introduction](#)

[Previously unnoticed/unaskable questions](#)

[Early compositionality and the Fundamental Construction Kit](#)

[What is compositionality?](#)

[*From sentences to pictures*](#)

[*Compositionality in computer science and AI*](#)

[*Extensional and intensional semantic compositionality*](#)

[*Biological compositionality*](#)

[*Sharp and fuzzy learning*](#)

[*Compositionality in evolution's construction kits*](#)

[*Construction kits for building mathematicians*](#)

[Turing on mathematical intuition vs ingenuity \(Added: 1 Nov 2018\)](#)

[Frege on compositionality](#)

[*Dealing with unshared linguistic context*](#)

[*Kant on geometric/topological reasoning*](#)

[Metaphysical Compositionality \(Added 19 Sep 2018\)](#)

[*Does Category Theory help?*](#)

[The importance of compositionality in Biology](#)

[Composition of process structures](#)

[Note on change-blindness](#)

[Alternative forms of compositionality](#)

[Waddington's "Epigenetic landscape"](#)

[Figure Wad](#)

[A meta-configured genome's use of compositionality](#)

[Figure Meta](#)

[Layered genome expression and use of construction kits](#)

[Non-quantitative semantic competences](#)

[Evolution of ancient mathematical abilities](#)

[Compositionality in natural language](#)

[Intensional and extensional compositionality](#)

[Reasoning about appearance changes caused by motion](#)

[Spider creativity](#)

[Compositionality in biological niches \(Added 2018/10/04\)](#)

[Cross-species/Symbiotic compositionality \(to be expanded\)](#)

[\(Added 19 Jul 2019\)](#)

[CONCLUSION](#)

[THANKS](#)

[APPENDIX A](#)

[A tutorial example illustrating linguistic compositionality](#)

[Figure Blocks](#)

[Figure Parse tree](#)

[REFERENCES](#)

[Some of my freely available online material related to this topic](#)

Keywords/phrases:

Compositionality in biological evolution;
Compositionality in individual development;
Compositionality in mechanisms produced by evolution;
Compositionality in biological information processing;
Compositionality in epigenesis;
Compositionality in biological niches;
Compositionality as a source of mathematical domains;
Discrete compositionality (as in logic, algebra, lego bricks)
 vs continuous compositionality (e.g. plasticine, water, rubber bands)
 vs hybrid (discrete+continuous) compositionality (e.g. meccano, Euclidean geometry)
Compositionality as a source of mathematical problems and competences;
Compositionality and creativity.
 Compositionality as a source of creativity in evolution;
Uses of mathematical compositionality in evolution and its products;
[Compositionality in fundamental and derived construction kits](#);
[The Meta-Morphogenesis project](#)
Metaphysical compositionality

Preface: Evolution uses increasingly abstract evolved and developed types

Most philosophies of learning and science assume there's a fixed type of world to be understood by a learner or a community of scientists. In contrast, biological evolution, in collaboration with its products, including social/cultural evolution, is highly creative metaphysically, constantly extending the mathematical and metaphysical diversity and complexity of the world to be understood, while constantly extending the types of learners with new powers. Its more advanced (recently evolved) learners have multi-stage genomes that produce new powers within each learner, suited to extending the powers produced by the earlier stages and their environments. This involves continually adding new, more abstract and general levels of compositionality during both evolutionary and developmental transitions and indirectly during cultural evolution. As explained [below](#), meta-configured genomes allow previous products of evolution to be combined in new ways within individuals, without altering the genome. This is related to Haeckel's suggestion that "Ontogeny recapitulates phylogeny", except that the recapitulation of cognitive capabilities is at levels of abstraction that allow enormous variation across generations. https://en.wikipedia.org/wiki/Recapitulation_theory

Introduction

Compositionality occurs when properties of a complex structure depend, in systematic ways, on properties of and relations between components of the structure. E.g. the thought expressed or scene depicted by a complex sentence or picture depends on what is expressed or depicted by the parts and how those parts are related in the sentence or picture. Some properties (e.g. truth of a statement, or misrepresentation in a picture) may depend also on how parts, and the whole, are related to something external, e.g. parts of the world.

Compositionality also affects *modal* relations, such as consistency, implication and incompatibility, between different structures with semantic contents, and is thereby a source of mathematical necessity and impossibility. Recognition and use of such necessities and impossibilities underlies human mathematical cognition (as noted by Kant ([1781](#))), as well as engineering creativity and scientific explanation.

Unlike neurally inspired statistics-based AI, which deals in probabilities, these modal features can be found in logic-based AI, but only in limited forms that do not, for example, explain the possibility of ancient mathematical discoveries in topology and geometry, or corresponding forms of spatial intelligence in humans and other animals, illustrated below.

Following Gottlob Frege (whose views are discussed further [below](#)) compositional concepts have been applied to linguistic utterances, and larger linguistic structures, e.g. paragraphs, or stories. One of Frege's main achievements was pointing out that quantifiers (For all..., For some, etc.) can be construed as higher order functions that take functions as arguments (i.e. predicates construed as functions from objects to truth-values).

Since the 19th Century, compositionality has been investigated in a wide range of complex phenomena, including, more recently, computer programs, programming languages, mathematical formalisms, proofs, pictures, maps, diagrams, musical compositions, thoughts, percepts, intentions, and other mental contents, though the word "compositional" is not always used as a label for the phenomena.

Thanks to the work of Frege, Tarski and others, it is well known that compositionality pervades human languages and logical/mathematical formalisms, where semantic content and semantic relationships are propagated "upward" in a systematic way from non-logical components through logical operations to larger structures. As we'll see, there are additional varieties of "content-construction", in many of which both non-linguistic information structures and non-linguistic background information have important roles. (This essay will not consider possible exceptions such as poetry and other "expressive" uses of language.) Compositionality is also a major source, perhaps the only source, of mathematics, illustrated below.

Compositionality is also a feature of the physical universe. On this planet, mainly as a result of biological evolution's use of the variety of potential resources provided by the physical universe, especially chemistry, there has been, over billions of years, steadily increasing diversity and complexity in physical, chemical, biological, and more recently mental and human-engineered structures and relationships, constantly generating new, more complex, forms of compositionality.

Even if, according to the 2nd law of thermodynamics the diversity will eventually be consumed in a global "heat death", on this planet we are still on a trajectory away from that state, fuelled by solar energy and locally stored energy.

The process has so far mainly been driven by evolution's constant production of new life-components, new construction-kits, new life-forms, new behaviours, and new products of those behaviours, including complex physical constructions such as birds' nests and termite cathedrals, with kinds of compositionality of structure and function that could not have been produced purely by physical/chemical processes without any use of information-processing mechanisms that are at the core of life.

For centuries, however, human decision-making and products of human engineering have been accelerating that process, through production of new technologies and new products of those technologies.

Developments in information processing technology (starting before computers, with machines like Jacquard looms, music boxes, uses of punched cards in the 19th Century), have now increasingly accelerated the rate of production of new structures (partly through more sophisticated control of other forms of machinery). General systems theory and cybernetic theory were attempts to characterise that process, but lacked computational sophistication, to the extent that they used numerical rather than more general computational models, e.g. manipulating grammars, symbolic data-structures, molecular mechanisms, etc.

Compositionality in the universe explains why the combination of mathematical representation and reasoning plays a central role in scientific explanations. However, sciences that focus only on numerical features, e.g. relationships between measures, or frequencies of co-occurrences, as is common in the social and behavioural sciences ([Andreski,1974](#)), miss out on the deepest [structural](#) features of the universe.

Newtonian physics with only interacting point-masses with elasticity, gravitational attraction, position, velocity, etc. cannot account for interacting enduring rigid, or topology-preserving structures, like molecules, whereas quantum mechanics, despite its notorious nondeterminism, can, as explained in [Schrödinger\(1944\)](#). (I suspect Newton was aware of this limitation in his otherwise outstandingly successful theories, and that may explain why he was so interested in Alchemy.)

It seems not to have been widely acknowledged that compositionality is involved in many forms of biological information processing, including epigenesis, i.e. gene expression. Using ideas developed with Jackie Chappell, I'll try to show that we can improve on Waddington's idea of an "Epigenetic landscape", summarised [below](#), and related ideas developed since then, by pointing out that the mechanisms of epigenesis (gene expression) also make use of (upward and downward!) compositionality. As a label for this idea I talk about "*Meta-configured genomes*".

"Downward compositionality" in programming languages, has been used since the late 1960s in "inheritance" mechanisms deployed in misleadingly named Object Oriented Programming and related techniques in computer science and software engineering.

Previously unnoticed/unaskable questions

These relatively recently developed concepts can be used to formulate some previously unaskable (or difficult to express) questions, including questions about how biological compositional mechanisms function in various aspects of reproduction, growth, development, learning, cognition, motive formation, plan formation, selection and execution, collaborative actions, and, in some species, culture formation, and other biological phenomena. We can also formulate new questions about whether and how this functionality can be modelled or replicated in symbolic or neural or hybrid AI systems.

Can they be implemented using current programming languages and the discrete computational machines on which they run, or do we need new fundamental forms of computation, e.g. chemistry-based computation, which Turing seems to have been investigating shortly before his death [Turing\(1952\)](#), perhaps partly inspired by the ideas in [Schrödinger\(1944\)](#) and [Craig\(1943\)](#)?

Some researchers may wish to argue that some form of quantum computation is required for intelligent life since there is already evidence that it is used in organisms for other purposes, e.g. [Seth Lloyd\(2012\)](#).

My aim in this paper is not to answer all the questions but to draw attention to some that seem not to have been noticed by most researchers. I'll also make some suggestions for further research, and mention some of the history of these ideas that is not widely known.

My own (initially unwitting!) work on this topic began while writing my 1962 DPhil thesis defending Kant's philosophy of mathematics [Sloman\(1962\)](#), but the ideas were expanded dramatically later, initially when Max Clowes ([Sloman-Clowes-1984](#)) introduced me to AI in 1969, including his own work on natural and artificial vision, and later in discussions about epigenesis, with Jackie Chappell, since 2004, as reported in [Chappell & Sloman\(2007\)](#) (expanding our 2005 IJCAI paper into an invited journal article).

Some key features of our theory are explained below, referred to as "[The Meta-Configured Genome](#)".

Being invited to contribute to the Elsevier volume commemorating Alan Turing's life and impact [Cooper & van Leeuwen \(Eds.\) 2013](#) prompted the question: why was Turing writing a major paper on chemistry-based morphogenesis shortly before he died (included in the volume). My conjectured answer was that he was thinking about what I labelled "The Meta-Morphogenesis Project" [Sloman \(ongoing\)](#) -- a study of forms and mechanisms of evolution, and ways in which products of evolution can change the scope of evolution. A subset of that is a new approach to developmental psychology.

Early compositionality and the Fundamental Construction Kit

Biological evolution starts with the Fundamental Construction Kit (FCK) provided by the physical universe, or to be more precise, the FCK augmented by a variety of physical products, e.g. galaxies, solar systems, planets, and a large collection of intermingled, interacting, types of matter, undergoing various forces and types of interaction, including chemical interaction, on, in, and around planets, continuously bombarded by various forms of radiation, especially thermal radiation from the sun, on this planet. Exactly how evolution starts (e.g. what the first evolvable forms were) is unknown. Once begun, evolution, augments the FCK with an increasing variety of increasingly complex Derived Construction Kits (DCKs) of many kinds (including construction kits for building construction kits), thereby repeatedly extending evolution's powers, and enriching the forms of compositionality to be described in an explanatory theory.

(Could Category Theory be relevant here??)

Corresponding to Frege's distinction between Sense (Sinn) and Reference (Bedeutung) (often indicated by other labels, such as intension vs. extension, connotation vs. denotation) there are two kinds of linguistic/conceptual compositionality that can coexist: intensional and extensional compositionality, as explained below. This has consequences for philosophy of biology, logic and mathematics, as well as philosophy of mind and philosophy of language, the main context in which compositionality has been discussed.

Facts about compositionality discussed below need to be recognized both by designers of AI language understanding systems and, more generally, researchers attempting to use AI techniques to model information processing in intelligent machines, including robots interacting with complex changing spatial configurations. (Compare squirrels and weaver birds.)

I suspect that compositional information processing mechanisms required for such purposes were at the root of much (informal) mathematical discovery long before the development of modern mathematics and meta-mathematics. Those mechanisms have not yet been replicated in AI systems. It's not clear that they can be replicated on digital computers. Perhaps that is why Turing was working on chemistry-based morphogenesis/compositionality shortly before he died?

What is compositionality?

There is a familiar notion of compositionality used in theoretical linguistics, formal semantics, mathematics and meta-mathematics (where it all started, building on work by Peano, Frege, Russell, Hilbert, Tarski, Robinson, and others). There are likely to be other notions using the same label for different aspects of language, or cognition, or some other activity, but I have not attempted a complete survey. Some examples of linguistic compositionality and their use in a "toy" demonstration Pop-11 program that accepts simple commands and questions are shown in [Appendix A](#), inspired by the far more complex computational model reported in Terry Winograd's MIT thesis ([Winograd \(1972\)](#)).

What I've referred to as a "familiar" notion of compositionality assumes that there are complex entities that share some common structure but vary in the components of the structure, and each combination of structure plus components produces some (resultant) attributes (denotations, functions, uses, ...). The attributes of complex structures are determined in systematic ways by interactions between

- the fixed/invariant structure common to many instances, and
- properties of the components of the structures that can vary, including changing relations between components, e.g. shape changes of complex objects.

Examples with more than one level of compositionality include sentences in this document (like this one!) whose semantic contents depend on various grammatical structures whose abstract semantic properties are instantiated by being combined with lower level (more concrete) semantic contents determined by lower level combinations of structure and content, as exemplified in the way the semantic content of this sentence depends on its various levels of syntactic structure, and the ways in which individual words or part-words (e.g. tense or number indicators) contribute to the semantic contents of larger parts, which in turn contribute to semantic contents of still larger parts. (That example was deliberately tortuous to illustrate the point it makes.)

The simpler kinds of linguistic compositionality are often illustrated by means of parse trees, showing grammatical structures, annotated with indications of what the various sub-trees refer to. A toy example is provided by the program presented in [Appendix A](#). In the example, a question is asked about a scene containing blocks. The program derives a *parse tree* for the question, showing its grammatical/compositional structure, then builds a *semantic interpretation*, another compositional structure, derived from the parse tree. It then answers the question by using the question structure to drive interrogation of a previously provided database of information about the scene. In a physical robot, the last step might be replaced by a perception system getting information from the physical environment, unless the information had previously been acquired, in

which case the answer could be derived from an internal memory structure.

Such examples, and their underlying theories are far too simple to account for the many functions of human languages, in part because they usually do not include a comprehensive account of functions of languages (e.g. their uses in thinking and reasoning, without any communication being involved), and they mostly (all?) ignore the evolutionary precursors of human languages that are required for many aspects of non-human intelligence.

The first such precursors must have had purely internal (non-communicative) semantic functions, e.g. expressing percepts, intentions, hypotheses, plans, beliefs, etc. A cat that jumps onto a dinner table after humans have left the room may have a belief about the availability of food, e.g. left over scraps. But there's no reason to think the cat's beliefs are exactly translatable into any human language. Likewise a weaver bird intending to make a new knot and then carrying out the intention will have a collection of beliefs and intentions and a changing set of percepts used for controlling its actions. None of the contents may be exactly translatable into English.

These *constructive* compositionality-based processes are different from *probabilistic* learning-based processes, where each process instantiates a transition pattern learnt from previously encountered examples, where the choice between alternative transitions depends on the frequency of previous associations. Such a theory could not explain how you can read and understand an essay like this, in which most sentences that are more than a few words long have not been encountered previously.

In this paper I shall not attempt to explain why statistics-based probabilistic learning fails to match the abilities of brains to produce new ideas, or new elaborate theories, stories, jokes, poems, pictures, musical compositions, or creative excuses for behaviours likely to lead to punishment! Neither can such learning explain how a squirrel is able to work out a way to defeat a new "squirrel proof" bird-feeder.


Moreover, reasoning from the fact that a previously acquired abstract pattern with mathematical properties (e.g. impossibility or necessity of some combination of features) is applicable to a newly observed phenomenon is a very different process from extrapolating observed co-occurrences using statistical reasoning. A varied collection of examples of perception and reasoning about impossibilities and necessary connections can be found in [Sloman\(impossible\)](#). Also in Euclid's *Elements*!

From sentences to pictures

In the 1960s, some AI researchers, e.g. Max Clowes, [1971](#), [1973](#), David Huffman [1971](#) and others (e.g. see [Kanef\(Ed\) 1970](#)) extended concepts of "syntax" and "semantics" to images and diagrams, and proposed computational mechanisms for assigning semantics to 2-D depictions of 3-D structures, maps, circuit diagrams, etc. Where a 2-D structure depicts a 3-D scene, working out what the scene is, is a non-trivial process, since assigning compositional semantics to complex images often involves solving constraint problems in order to resolve local ambiguities. That is also a feature of biological perceptual processes. Here's an example:

Here are nine blocks above a surface, shown in different configurations, in some of which one more of the blocks is suspended above (or floats above) another block. How many other configurations are possible? How else could they be arranged? What sequences of block moves could transform one of the arrangements into another arrangement?


Are all the depicted arrangements spatially possible? If there is an impossibility what sort of visual system can detect it? Would you expect a two week old human infant, or a squirrel, to be able to detect the impossibility? If not, what mechanisms do they lack?



See also

http://psylux.psych.tu-dresden.de/i1/kaw/diverses%20Material/www.illusionworks.com/html/art_of_reutersvard.html

The above examples were all inspired by a picture drawn by Oscar Reutersvard some time before the Penrose triangle and Escher's pictures of impossible scenes.



The ability to detect impossibilities depends on abilities to reason about structural relationships, not the ability to generalise from masses of examples. Could a deep learning system be taught to recognize spatial impossibility in 2D images of 3D scenes? How would it represent (or understand) "impossible"?

In this case, and many of Escher's examples inspired by the Penrose triangle, the impossibilities arise from composition of spatial relationships like: *further*, *higher*, *more left* and *more right*, each of which is transitive and antisymmetric. How can a biological brain represent and reason about such relationships, and understand spatial impossibility and necessity?

Young children don't all see the impossibilities that are obvious to older humans. What has to change in their brains during development?

Jean Piaget investigated this question, but he lacked computational knowledge.

Most psychologists and neuroscientists don't understand the question.

Many current AI researchers seem to be equally ignorant: they think all spatial perception is a form of identification or classification. Contrast Gibson. Others think of perception as concerned with finding probabilities that various actions will produce certain consequences. More on that below. For examples of early computer vision work see [Ballard and Brown\(1983\)](#). Some challenges to all AI/neural/psychological theories so far are presented in [Sloman\(impossible\)](#).

Compositionality in computer science and AI

Parallel developments in Computer Science and AI, including use of specially designed programming languages, allowed semantic content to propagate not only "upward" from parts to wholes, but also "downward" in various forms of inheritance in "object oriented" programming ([OOP](#)), including programming languages with multiple inheritance, "Semantic Nets", and "Frame Systems" in AI (starting in the 1970s).

Biological compositionality

The Turing-inspired "Meta-Morphogenesis" (M-M) project, proposed in [Sloman\(2013\)](#), extended in [Sloman\(ongoing\)](#) generalises these ideas, suggesting that biological evolution anticipated (and in some ways went far beyond) recent ideas about compositionality. Evolution's [implicit](#) (or [unconscious](#)) creativity, based on types of compositionality discussed in this paper, exceeds every other known form of creativity!

For example, during development of an individual organism, biological functionality results from a combination of information flowing "down" from the genome, and flowing "up" from results of earlier interactions between genome and (internal or external) environment. Current ideas about learning and development have explored only a shallow subset of such interactions, e.g. using statistical learning.

In more complex cases (e.g. mammal development) a genome has components that are not "expressed" all at once, but in stages, where later stages can be transformed by products of earlier stages, as proposed in [Chappell & Sloman\(2007\)](#). (Compare [Karmiloff Smith \(1992\)](#).) This is explained in more detail [below](#). Human language development is a spectacular example.

The discovery of the structure of DNA announced in 1952, as summarised very briefly in this Science History Institute [article\(2017\)](#) was a remarkable example of compositionality of information-processing distributed across four (mainly three?) different scientific minds with different knowledge and competences, combining intellectual powers and newly acquired factual information to produce a world-shaking theory.


The DNA replication mechanisms they described (at least in outline) were also examples of chemical compositionality involved in reproductive processes whose implications scientists are still discovering. One of the features of those processes is that unlike the types of compositionality found in modern algebra and logic, and theoretical linguistics, the compositionality of the molecular processes essentially involved a mixture of discrete and continuous processes. The discrete processes created or removed chemical bonds between atoms. The continuous processes involved movements of complex molecules through space, including translation, folding, and twisting. Similar comments apply to compositionality in visual perception, especially perception of moving, interacting, structures, in complex machines, or dancers performing a ballet.

There seems to be a growing body of evidence and theory suggesting that related molecular processes *within* a synapse play a far more important role in animal cognition than hitherto realised, challenging popular theories of learning based on connections *between* synapses, summarised in [Trettenbrein\(2016\)](#). (Von Neumann anticipated such a development in his 1958 book, *The Computer and The Brain*, written while he was dying.)

Compositionality in evolution’s construction kits


Production of new species, and development of their individuals, both require use of the *fundamental* construction kit provided by physics and chemistry, and also increasingly many kinds of *derived construction kit*, discussed in more detail in [Sloman\(2017\)](#) and [Sloman\(kits\)](#), including construction kits for producing and modifying construction kits (meta-construction-kits??).

Figure FCK: The Fundamental Construction Kit



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

Figure 2 DCK: Derived Construction Kits



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

The ability to detect depictions of impossible 3D objects would require special "modality" detectors, e.g. for classifying thought of or depicted spatial structures as [impossible](#). Examples were developed by AI vision researchers, e.g. Max Clowes, around 1970). Evolution presumably produced such impossibility detection mechanisms, or mechanisms for building them, not because they are of mathematical interest but because they are practical importance, e.g. in ruling out futile construction ambitions. (There's a lot more to be said about this.)

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

Some biologists have noticed the need for these ideas, e.g. [Coates et al. 2014](#). (They refer to "toolkits". Their work triggered my thoughts about construction kits containing both tools and materials.)

Derived construction kits of various sorts are required (recursively) for production of increasingly complex forms of physical material (e.g. bones, muscles, sinews, nerves, skin, genetic materials), and use of those materials in different structures in plants, fungi, bacteria, etc.

Construction kits are also required (recursively) for information processing mechanisms, new varieties of information, new biological functions, and for multiple layers of increasingly complex derived construction kits required for continued evolution of complexity, or new kinds of individual development, using new kinds of compositionality.

Man-made construction kits, such as meccano, tinkertoys, lego, fischertechnik, origami, dress-making and knitting materials and tools, can all be viewed as construction kits with compositionality. In all cases there are component materials, actions and tools that can be used to produce increasingly complex structures, whose causal powers and constraints emerge

systematically from properties and powers of parts and the relations between them.

Continuous and continuously deformable materials such as mud, plasticine, dress-making materials, and sand (viewed as stirrable, pourable, etc.) all can be seen as providing extensions to the ideas of discrete compositionality. Origami is an intermediate case, in which pre-existing discrete items are not brought together, but are created by folding paper, to produce new 2D portions of the original 2D sheet, and also 3D structures and relationships.

Like chemistry, many "artificial" construction kits include a mixture of discrete and continuous compositionality. For example in meccano a pair of perforated strips can be linked by a nut and bolt passing through a hole in each part: a discrete change. But the if the nut and bolt are not fully tightened continuous variation of the angle between the two parts remains possible. A child playing with meccano will pick up abilities to reason about different combinations of parts, including understanding which new possibilities and impossibilities are generated.

Spatial movements and changes of viewpoint can also be understood as generating networks of continuous variation linked by discrete changes, e.g. seeing the visible gap between two objects grow smaller then disappear as the objects come into contact, or one part passes behind the other. If the motion continues, the remaining visible portion can grow smaller until it is completely out of sight, a discrete change. Many intelligent animals need to be able to make use of such information in selecting and controlling their actions. Humans are also able to reflect on what they are doing, and in some cases to communicate what they have learnt to others. I suspect these facts are part of the original source of ancient discoveries in Euclidean geometry, starting long before any records were kept.

Euclidean geometry as a whole, and various subsets of it can also be seen as illustrating rich and deep forms of compositionality where new properties, relations, and constraints arise in larger topological/geometric structures through the interaction of properties and relations of parts. We'll also see later that biological evolution makes use of varieties of compositionality in which new more complex structures and mechanisms support new varieties of functionality, often implicitly making use of mathematical facts.

Evolution's uses of these new "discoveries" suggests a view of evolution as a *blind mathematician*. Dawkins' notion of a *blind watchmaker* is a special case of this.

Despite all these grand claims, it seems that very little is understood about the evolution of most of the phenomena mentioned here.

Construction kits for building mathematicians

A challenging test case for these ideas is to identify the (multi-layered) products of evolution that enabled ancient mathematical discoveries to be made, e.g. by Archimedes and many others before and after him, long before use of modern logic and the relatively recent Cartesian interpretation of geometry in terms of numbers and numerical operations. So far, neither neuroscience nor AI seems able to explain or model those ancient achievements, or their precursors in other animals with spatial intelligence.

It is clear that the genome alone does not suffice for production of human mathematical competences: there had to be, in the environment of the baby/young Archimedes, products not only of physical construction kits, but also products of earlier products of the human genome,

including not only immediate and remote human ancestors (and their near and remote relatives) but other conspecifics.

In short, each stage of expression of the human genome in an individual benefits from many products of the expression of the genome in previous generations and in parallel branches, including toys and other playthings produced in the local culture that help to shape the development of young human minds, as well as benefitting from products of earlier stages of genome expression.

One fairly obvious example is the relationship between the "babbling" stage of development which helps to produce competences required for producing sequences of word-fragments suited to the language spoken in the infant's environment, and a later "grammar+semantics" stage where, after word competences have been developed, phrase and sentence competences begin to be developed -- partly driven by the genome and partly driven by the current linguistic environment.

Evidence from deaf children in Nicaragua indicates that the genome is more important than the environment if there are enough conspecifics developing in parallel, as made clear in this short video:

<https://www.youtube.com/watch?v=pjtioIFuNf8>

What is not clear from the video or most reports on language development, is that [internal](#) languages must be developing too, partly under the control of the genome and partly under the control of the environment, including languages for expressing perceptual contents, intentions, plans, reasoning processes, memories, and many more. This will be true, to a lesser extent, of other intelligent species who sophisticated language use is completely internal.

The products of evolution involved in mathematical development include rich [internal](#) (not necessarily linear, or discrete) languages used for percepts, intentions, beliefs, questions, control of actions, and other functions. For internal information structures derived from visual perception the form of representation may be *analogical* (not *analog*) in the sense of Ch 8 of [Sloman\(1978\)](#), which contrasts Fregean vs Analogical forms of compositionality. It is possible that at least the internal languages of pre-verbal humans are partly shared with other species. (As far as I know, the distinguished mathematician, David Mumford, is the only other person to have appreciated the need for such internal languages: see his online [Blog](#).) Human "external" languages must have evolved later, but are not (and cannot be) fully encoded in genomes, since they differ wildly around the planet.

I am not claiming that there is an innate language of thought of the sort proposed in [Fodor\(1975\)](#). For a recent discussion of that issue see [Beck\(2018\)](#). There are also useful related discussions in [Hampton\(2017\)](#).

Human brains, like the brains of ancient mathematicians, include mechanisms (collections of interacting [virtual machines](#) -- discussed in [Sloman\(2013,ff\)](#)), with sub-mechanisms that are closely related to mechanisms in brains of many other intelligent species, e.g. squirrels, crows, and elephants, and pre-verbal human toddlers. It seems to me that those abilities cannot be explained by current psychology or neuroscience, or fashionable "Deep Learning" systems. We may need revolutionary ideas, especially about natural forms of information processing, including evolved varieties of compositionality.

Turing on mathematical intuition vs ingenuity

Added: 1 Nov 2018

One of the arguments for this was unintentionally prophetically provided by Alan Turing in 1938, when he made a distinction between mathematical intuition and mathematical ingenuity, claiming that only the latter can be (re-)implemented on digital computers. At that stage he left open the question: what sorts of mechanisms could provide an implementation for forms of representation and reasoning needed for developing and using mathematical intuitions. This is still a major unsolved research problem for AI and neuroscience. That topic is discussed further in this summary and discussion of Turing's 1938 statements: [Sloman\(on Turing1938\)](#).

Sharp and fuzzy learning

Compare: use of mathematical structures and processes produced by composition to produce sharp conclusions that don't need to mention probabilities of correctness.....

with use of "soft" information (statistical correlations) to produce probabilities rather than definite conclusions. Why can't linked solid impermeable rings become unlinked? vs how much force is needed to separate them by breaking one?

Frege on compositionality

Although Frege is widely acknowledged as having contributed deep new ideas concerning compositionality in human languages and formal languages, [Pelletier\(2001\)](#) documents a considerable amount of uncertainty and disagreement among Frege scholars about exactly what Frege did and did not say, or think, at various times. Although much of my thinking was stimulated by studying Frege (for my DPhil), I felt then (and now) that he had missed an important feature of compositionality outside mathematics and logic, insofar as he seemed not to pay enough attention to the fact that in many cases the result of composition (e.g. what is referred to, or denoted) depends partly on how the world happens to be, as explained below, which is not generally the case in mathematics, his main concern. Anyhow, I hope the claims made here are intelligible and discussable in their own right, independently of what Frege wrote or thought.

Some of the ideas were presented in [Sloman\(1962\)](#), acknowledging Frege's insight that semantic complexity often arises out of application of functions (including higher-order functions, such as quantifiers) to arguments of various sorts. I thought he did not deal adequately with the fact that many of his functions take the state of a relevant part of the universe, as an implicit additional argument, as explained in [Sloman\(1965\)](#). This is obvious for examples like,

"The current president of the USA",

"The current president of the USA is a mathematician",

"The distance between the centre of the earth and the centre of the sun on 9/9/1999".

I (unfortunately?) used the label "rogator" for such a "world-dependent" function, derived from the Latin *rogare* "to ask", because rogators, unlike pure mathematical functions, have to "ask the world a question", in order to get a result, though in special cases the semantic content of the question determines the result independently of how the world is e.g. "Everyone over one meter tall is over two centimetres tall", whose truth is determined by the defining relation between "meter" and "centimeter".

One consequence of these ideas is that in some cases the referent of a complex linguistic expression, or part of the expression is not determined by the linguistic context of use alone: the state of some part of the non-linguistic universe also plays a role. A proud parent using the

expression "my youngest child" on two occasions a year apart may refer to two different individuals. (Is this something category theoretic models of language can cope with?)

Dealing with unshared linguistic context

Linguistic communication (and thinking) includes far more complex and subtle cases where what is expressed by a particular word, phrase, or sentence depends not merely on aspects of the shared environment, but on aspects of a speaker or hearer, or reader's current mental content. For example, if I don't know that you have acquired a new car, I'll mistakenly interpret your offer "I'll take you in my car" as referring to the car I am already familiar with, and when you refer following some social interaction to a piece of very rude behaviour, our different social values may lead me to misidentify the action you are referring to. Good novelists and playwrights understand all these matters intuitively and use them in setting up scenarios in their novels. A clever, manipulative, politician could deliberately choose a context-sensitive mode of expression knowing that different parts of the audience will interpret it differently, and approve for quite different reasons.

More complex cases where simple theories of compositionality will break down include cases where what is written or said implicitly depends on a rich background in the speaker's mind that the speaker is incapable of communicating, but has to rely on a creative response in readers or hearers, for example in poetry, or an early formulation of a new scientific theory. Such unobservable individual differences can undermine (naive?) empirical research on how language works. Compare [Hampton\(2017\)](#).

This is a huge topic, involving many different connections between linguistic and non-linguistic competences and mechanisms. By comparison with required explanatory theories, much psychological and psycholinguistic research seems to me still to be in a "Ptolemaic" phase: trying, like ancient astronomers, to find patterns in data, while lacking adequate explanatory theory construction capabilities. Compare [Andreski,1974](#). These problems also beset data-based AI research on linguistic competences, vision, reasoning, etc.

AI desperately needs deep new explanatory theories about products of biological evolution -- theories that go well beyond data available from observations and experiments so far, as has repeatedly been the case in physics and chemistry, at least in the last few hundred years.

Kant on geometric/topological reasoning

There are more problematic cases noticed by Kant, for which a deep understanding of some generic aspect of the world suffices to determine a truth-value, without the need to check any particular cases, e.g. "It is impossible (in a Euclidean space) for three planar surfaces to completely enclose a finite 3D volume": an example of the sort of mathematical knowledge that Kant described as *a priori* (non-empirical), *synthetic*, not derivable from definitions alone using pure logic, and *necessarily true*. His claims seem to have been rejected by most philosophers of mathematics, but from my own experience of doing mathematics (geometry especially), I am sure he was right, as may eventually be demonstrated by building a "baby" robot, that develops the ability to make such discoveries. For that we need to unlock some still hidden secrets of the human genome, by expanding the ideas presented [below](#) about the "Meta-configured genome".

At present, I know of no plausible account of the brain mechanisms that allowed ancient mathematicians to make discoveries about necessary connections and impossibilities involving space; mechanisms that still allow young mathematicians to repeat some of those discoveries, and occasionally make new ones.

As far as I know, those who propose neural models of mind ignore this sort of requirement, though the requirement was understood by the developmental psychologist Jean Piaget: his last two books were about *Possibility* and *Necessity*, though as far as I know, he never mastered AI or programming, and looked for explanatory mechanisms in the wrong places.

Current machines programmed with the abilities to discover and prove theorems in geometry, starting from [Gelernter 1964](#) all, as far as I know, make use of modern logical and arithmetical apparatus not used by ancient mathematicians. For example, the earliest geometers could not have known Euclid's *Elements*, and did not make use of the Cartesian mapping from geometry to arithmetic and algebra, based on numbers as coordinates. Such machines can produce some impressive results, but clearly not using the modes of spatial reasoning available to the ancients, or to a child, like this pre-verbal toddler (about 17.5 months)

<http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/small-pencil-vid.webm>

and I am not aware of any AI system with the spatial intelligence of a squirrel, weaver bird, elephant, orangutan, etc.

I'll try to show elsewhere that animal spatial intelligence makes essential use of spatial structures and processes in what James Gibson called "the optic array" (i.e. not patterns on the retina, but patterns in the light cone converging on the eye -- the subject matter of much painting and graphic art). Those structures and processes are related in complex ways not only to the structures and processes that exist in the environment independently of the viewer, but also to systematic changes in information available caused by interactions between:

- viewer motions,
- static spatial relations in the environment, and
- motion of some parts of the environment, including
 - Objects moving in and out of view, or partly in and out of view;
 - Objects varying their relative distance from the viewer, and spatial relations with one another;
 - objects absorbing other objects, including eating.

Often absolute measurements don't provide the most useful detectable types of *change* (e.g. increasing, decreasing, decreasing with deceleration, etc.).

These features change not only according to the distance from the viewer, the motions of viewer, object, and intervening objects, but also changing orientations of seen objects altering aspect-ratios, changes of visible projected gap sizes between visible edges, changes of projected angle sizes, and changes in perceived orientations of visible edges or gaps. Current AI systems would normally assign numerical values to such changes, modified by probabilities. I suspect there are far more powerful, more rapidly computed, *qualitative* descriptions (e.g. increasing, decreasing, accelerating, increasing more than or less than something else, etc.)

Another powerful ability, at least in humans, and I suspect also in some highly intelligent non-human animals, is the ability to detect and reason about spatial impossibilities and necessary connections, which can save a great deal of unnecessary learning time. A varied, but in no way complete or systematic, collection of examples is available here, including links to additional examples:

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

Learning via huge numbers of observed correlations between such changes and muscular signals etc. is possible in principle, but I suspect (along with Kant) that evolution produced far more subtle and powerful spatial reasoning mechanisms that use necessary geometric and topological connections rather than statistical evidence. The ability of properly education young humans to work out ways of doing such reasoning in novel situations is strong evidence that brains have compositional reasoning mechanisms that have not yet been identified. But I do not claim that any of the mechanisms is used infallibly in the context of a working brain. Even great mathematicians can make mistakes, that are later detected and dealt with, as shown by [Lakatos\(1976\)](#).

As with linguistic competences these powerful mathematical mechanisms are not available at birth, and are not fully specified in the genome, but are created through processes of development that, as far as I know, nobody understands at present, using brain mechanisms that nobody understands (perhaps sub-neural chemical information processing mechanisms mentioned above [Trettenbrein\(2016\)](#)). This may be an instance of the operation of the Meta-Configured genome discussed [below](#).

Metaphysical Compositionality(Added 19 Sep 2018)

This sub-section is a place-holder. In many of the examples of compositionality in biology and engineering, especially computing, new *kinds* of entity come into existence for the first time (at least on this planet), including new kinds of virtual machinery, about which we have known only since about the mid-20th century, although long before humans discovered virtual machines, evolution discovered and used their powers, implemented on neural and sub-neural mechanisms. In organisms without brains or nervous systems, implementations were based only on chemistry (still in use in brains, e.g. in sub-neural chemistry [Grant, 2018](#)).

The mechanisms that made ancient forms of mathematical cognition possible must have been relatively late developments, but long before that evolution was implicitly using mathematical discoveries in its products, e.g. homeostatic control mechanisms (using negative feedback).

It is common to think that virtual machines are in some sense "merely" complex physical machines thought about in a useful short-hand way. That dismissive view is based on a failure to understand the extent to which virtual machines can introduce new *types of entity*, whose description requires new concepts that cannot be *defined* in terms of the concepts of physics, even though all their actual instances are *implemented* in the physical world ([Sloman\(2013,ff\)](#)). Since those concepts are capable of having new instances based on newly discovered biological or engineering developments, the concepts cannot be regarded as defined in terms of instantiations used. (This point is not widely understood, even by computer scientists!)

This often involves use of new kinds of *causation* and *causal-powers*. For example, a network of virtual machines can be running on a collection of computers whose physical components change from time to time e.g. because of repairs, or because of network-re-routing mechanisms, or because memory management systems in a single computer frequently re-locate portions of virtual machinery in the physical memory. The causal powers of the running virtual machine or machines are preserved through such changes. (Ancient philosophers noticed similar facts about complex physical entities, such as animals, plants and rivers.)

Perhaps the deepest reason why virtual machinery introduces metaphysical novelty is that a virtual machine running on one or more physical computers can be closely monitoring, reasoning about, and controlling a complex *external* system, for example a chemical plant, or a car assembly line in

a factory. These are not physical control relationships (as in a centrifugal governor), because they depend crucially on semantic structures and processes in the controller. As some new entity or new state in the controlled system is detected, this can trigger planning or problem solving or some remedial activity in the monitoring virtual machine. So new information can cause the contents of the virtual machine to change, including creation of new *semantic* references to items in the environment (e.g. hypothesised causes). Such intelligent control systems cannot be achieved simply by using increasingly complex physical control loops reactively responding to sensed physical values or mathematical functions computed from those values by adjusting forces, thresholds, angles, and other measurable components of physical states and processes. [Compare the ideas in "Cybernetics"(Wiener and Ashby), "General Systems Theory"(Bertalanffy) and "Perceptual Control Theory"(Powers).]

Referring is not a physical state: statements describing some entity *A* as *referring* to objects in its environment cannot be equivalent to a collection, or disjunction of collections, of statements describing the physical components of *A*, or even to a disjunction of statements describing *A* and its environment. Rather a state of referring is important because of its potential role in a variety of more complex states and processes, e.g. exercising control, avoiding, making plans, etc. The concepts used in describing these states and processes refer to components of a complex abstract structure whose states, processes, powers and effects are part of layer of networked causal relations whose descriptions require concepts that are not definable in terms of the concepts of physics, e.g. "detecting a near failure", "recognising a flaw in an argument" and many more descriptions applicable to complex control systems. (This requires more detailed discussion and references. Compare the contrast between Jane Austen's and Claude Shannon's concepts of "information" in [Sloman\(2013-2018\)](#).)

For similar reasons, statements about virtual machine components referring to other virtual machine components in the same or a larger virtual machine (e.g. simulated attackers as in [Kennedy\(2003\)](#)) are not equivalent to any collection of statements about physical structures, processes and causal interactions in the machine on which the referring virtual machine runs.

In that sense human engineers, especially software engineers, are often metaphysically creative. But the discussion of evolution in the rest of this document shows that biological evolution has been even more creative for a much longer time -- including forms of creativity required to produce new kinds of virtual machinery running on animal brains and other forms of control infrastructure.

This can be expressed as a claim that biological evolution discovered the power of control systems implemented in virtual machinery, running on complex biochemical machinery, long before human engineers discovered the power of virtual machines running on digital electronic machinery. Perhaps that's what Turing was thinking about when he wrote in [Turing\(1950\)](#): "*In the nervous system chemical phenomena are at least as important as electrical*". By then he may already have been writing, or at least thinking about the subject matter of, his paper on chemical morphogenesis, published only two years later [Turing\(1952\)](#).

Future research on this topic could explore the space of possible designs for control systems (brains? minds?) that can be implemented in mechanisms derivable from the fundamental construction kit provided by our physical universe, and various subsets of that space, e.g. those that can evolve without the use of any derived construction kits, and those that depend on various classes of derived construction kit. Biological evolution alone could not have produced our current networks of computers and the varieties of software (operating systems, compilers, applications,

internet and world-wide web tools, etc.), since they do not (yet!) include the required reproductive machinery.

In that sense human bodies and human capabilities can be seen as components of a very complex and changing type of evolved/derived construction kit that adds to the metaphysical possibilities supported by the physical universe. This needs an even more ambitious investigation into alternative possible universes and the types of mind that they can support, extending the discussion of the space of possible minds in [Sloman\(1984\)](#). Some of the requirements for such mechanisms are discussed in [Sloman\(2006--\)](#)

A future extension of this paper could include investigation of the relative scope of
-- virtual machinery implemented using digital circuitry, and
-- virtual machinery based on sub-neural chemistry,
expanding developments already reported in [Trettenbrein\(2016\)](#), [Grant\(2018\)](#).

I believe these ideas provide new, more complex, examples illustrating Wilson's ideas about "grounding as metaphysical causation", in [Wilson\(2017\)](#).

Does Category Theory help?

The symposium for which this paper was written emphasised both Compositionality and Category Theory. So far I have not noticed a role for category theory in the topics discussed here. Perhaps I have failed to understand the full power of category theory. (It was not included in any of the courses I took before switching from mathematics to philosophy.) Is it possible that category theory needs to be generalised in some way to accommodate multiple levels of interaction of structures, functions and forms of information found in biology, as investigated in the Meta-Morphogenesis project, of which this paper is a part?

David Corfield has written about the demise of geometry and whether it can be revived in [David Corfield \(2017\)](#). His questions are very different from mine. I ask about the roles of geometry, and geometrical information in biological evolution and development of intelligent animals. He asks more philosophically focused questions:

"I shall argue in this chapter that our best hope in reviving a 1920s-style philosophy of geometry lies in following what has been happening at the cutting edge of mathematical geometry over the past few decades, and that while this may appear a daunting prospect, we do now have ready to hand a means to catch up rapidly. These means are provided by what is known as cohesive homotopy type theory."

In contrast, my questions are about the roles of spatial cognition in organisms and in the operation of biological evolution, and how these biological origins were able to support ancient mathematical discoveries such as those reported in Euclid's *Elements* and still used world-wide by scientists, engineers, architects, etc. We are not in disagreement, merely asking different questions. I wonder if there is some deep overlap? Would Dana Scott's ideas on geometry without points, in [Scott\(2014\)](#), favour either? Does the claim that the concept of natural number depends on the concept of 1-1 correspondence support including arithmetic under geometry, if 1-1 correspondences are geometrical/spatial structures just as polygons, spirals, and projections of 3D structures onto 2D surfaces are.

[BACK TO CONTENTS](#)

The importance of compositionality in Biology

Compositionality has been discussed mainly in the contexts of linguistics, philosophy (especially philosophy of language), logic, mathematics, computer science, and sometimes physics. It appears not to have been generally noticed that compositionality pervades biology, including both "upward compositionality", the most widely known type, and "downward compositionality", whose role in epigenesis (gene expression during individual development) is discussed below. This raises important new research questions.

As a result of the work on foundations of mathematics and formal semantics by Frege, Tarski and others, followed by more recent work in theoretical linguistics, the role of compositionality in language, and therefore in linguistically expressed thoughts has been widely acknowledged. It is also central to the study of semantics of programming languages, and ways in which complex computational processes are composed of simpler ones. In all these cases semantic content is normally taken to propagate upward from simpler structures, through connecting operators (e.g. logical connectives) to larger structures. Since Frege, these ideas have been enriched by the addition of the notion of function-application, with predicates and relation symbols now construed as naming functions that can be applied to arguments.

This idea was generalised in the 1960s by AI researchers working on vision, e.g. Max Clowes, [1971](#), [1973](#), David Huffman [1971](#) and others. They extended concepts of "syntax" and "semantics" to diagrams and images of various sorts, and proposed computational mechanisms for assigning semantics to 2-D depictions of 3-D structures, a non-trivial process, since in many cases ambiguities of semantic content could be resolved only by solving constraint problems, as also happens in human languages where locally ambiguous components of a phrase, sentence, paragraph, or story can be disambiguated cooperatively, or interpreted in cooperation with other portions of the text. Visual disambiguation can also use background information, as illustrated by doodles and the examples presented in

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

The syntax/semantics distinction for images introduced to AI vision research [in the 1960s](#) was generalised in Chapter 9 of [Sloman\(1978\)](#) to explain how complex and messy "multi-layer" pictorial depictions such as this image




Fig 9-new In the revised online edition of [Sloman\(1978\)](#).

might be interpreted in a mixture of bottom-up and top-down processing dealing with structures in different domains, with different syntax and semantics, as (crudely) indicated below:




Fig 9.6 of Sloman(1978).

Some of the complexities and opportunities available in perception of processes and opportunities for processes (affordances) are presented in [Sloman\(2007-2017\)](#), discussed further in connection with mathematical competences [below](#).

There are many important cases of compositionality in biological evolution and its products. This paper draws attention to some of them and their relevance to deep unsolved problems in philosophy of mathematics, philosophy of mind, and mechanisms of biological evolution and individual development.

There are also deep challenges to AI theories of mental function, as well as to theories of biological cognition and neuroscience. More visual examples are available in [Sloman \(1983\)](#), and a collection of examples concerned with perception of possibility, impossibility and necessity here:

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

For a short, fairly abstract, summary of some of those ideas, including a demonstration of equivalence between some (but not all!) diagrammatic representations and logical representations see

<https://www.quora.com/What-are-Frames-in-Knowledge-Representation-in-Artificial-Intelligence>

An elementary introduction to Object Oriented Programming and its use of inheritance and compositionality, is available in an online Pop11 tutorial [Sloman OOP](#).

Insofar as Applied AI is part of software engineering, and overlaps with other forms of computer systems engineering, the points made here are potentially relevant to many forms of applied computer science, especially those in which complex new systems develop in a [spontaneous](#) but [constrained](#) manner at run-time, instead of all being specified in an initial design. Current AI "learning" systems and psychological or neural theories of learning and development typically address only a subset of the biological phenomena to be explained and replicated. Some of the more complex phenomena are described below.

In particular, the relevance of compositionality in genomes to mechanisms and results of individual development has not been widely noticed, especially by those who claim that Darwin's theory of evolution by natural selection based on random variation in traits suffices to explain everything. E.g. (in the words of Graham Bell ([Bell, 2008](#))):

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

But that does not explain what makes possible the options between which selections are made, options that are available in some parts of the universe (e.g. on our planet) but not others (e.g. the centre of the sun), and at different times on our planet: no evolutionary transitions available four billion years ago could have produced homo sapiens, or even squirrels, [directly](#).

Why not? Because over time natural selection produces not just novelty or increased complexity, but increasingly sophisticated [types](#) of design, where new layers of complexity build on old layers - as happens also with developing linguistic competences and developing computational designs for programming languages, information processing architectures, and practical applications of computing hardware and software. Some of the abstract design features required for mechanisms of squirrel or human intelligence did not yet exist on this planet four billion years ago. Evolution had not yet produced them. So the required forms of compositionality were not yet available, although precursors must have existed.

My impression is that many evolutionary biologists agree with Bell, quoted above, and do not notice the problem of explaining what makes evolution possible, or explaining how possible varieties and products of evolution change over time. Why was it impossible for humans to evolve directly from one of the earliest vertebrate species? What had to change to make evolution of humans possible?

The best known evolutionary theories (including Darwin's) offer only partial explanations. The work of Kirschner & Gerhart ([2005](#)) is one of the few exceptions I have encountered, though it is not clear that they recognised the requirements for [layered](#) construction kits mentioned below and discussed further in the Meta-Morphogenesis project (under continuing development in [Sloman\(ongoing\)](#)), or the different [kinds](#) of layering and compositionality, produced at different stages of evolution, discussed in more detail in [Sloman\(2017\)](#). I'll say more about these

requirements below.

The work of Paul Davies is also relevant [Davies \(2012\)](#), though I have not yet studied it sufficiently closely to be clear about the overlap. "Constructor Theory", being developed by David Deutsch and colleagues in Oxford <http://constructortheory.org/> is also relevant, though it seems to focus on only a subset of the issues, and makes assumptions about fundamental physics that I prefer to leave open, in case the biological phenomena turn out to require different fundamental physical mechanisms. If so, it will not be the first time that well known phenomena turn out to require replacement or enrichment of supposedly fundamental physical theories.

Note added 6 Jan 2022

This interview provides a useful summary of his ideas:

<https://www.youtube.com/watch?v=d8KCXgzqhSQ>

David Deutsch on 'Constructor Theory' 8 Dec 2018

Summary so far: sophisticated biological designs were as much beyond the scope of evolution in its earliest stages as sophisticated software engineering designs were beyond the reach of human programmers when they knew only about machine-code programming, or the kinds of data-structures and procedures found in early "higher level" languages, surveyed by Grzegorz Milka in:

<https://gregorias.github.io/2014/11/22/early-high-level-programming-languages.html>

I am suggesting that evolution had to discover new "higher level" design principles, supported by newly evolved "construction kits" via a bootstrapping process that is partly analogous to the way in which human programmers had to discover new design principles and support them with new, increasingly powerful, programming languages and tools. In particular, computer scientists and software engineers had to develop languages supporting various kinds of compositionality, including eventually "top down" compositionality (explained below). This has deep implications concerning the nature of life, and also requirements for tools to design and build human-like robots, a topic not discussed in this paper.

NOTE added 7 Jan 2022:

The claim that every possible form of computation can be implemented on a universal Turing machine (UTM) using a single tape with moving tape-head controlled by fixed rules is not true: there are many sophisticated self-reorganising forms of computation running on self-modifying virtual machines running on the internet while they constantly change the physical substructures used, which could not be implemented on a single Turing machine, e.g. a Zoom-based conference, airline reservation systems, internet-based multi-person games, banking systems, and many more.

Composition of process structures

Perception of objects in motion will require complex [processes](#) to be composed of simpler [processes](#). For example, perceiving someone walking involves perceiving [partly](#) independent motions of different body-parts, with different sub-processes changing the relationships between the torso, head, arms, legs, parts of arms, parts of legs, parts of parts, etc., and also changing relationships between those items and other static or moving items in the field of view, or recently perceived. The motions are only partly independent because of links between parts: e.g. the lower end of the upper part of a leg (the thigh) must coincide with the upper end of the lower part (shin/calf) of the same leg, and similar constraints correspond to all joints where two body parts are connected.

There are more complex and subtle constraints in a normal human environment: e.g. if no part of the body is in contact with any other surface, and there are no other force fields acting on the body apart from the Earth's gravitational pull, then the body cannot be static (relative to the Earth).

These constraints raise deep questions about how a brain or computer should represent a perceived process that lasts for a period of time, when earlier parts of the process are no longer visible -- they are in the past! Different problems arise when a complex static scene is surveyed by looking at different parts of the scene -- using a sequence of fixations, and when fixating a continuously moving blob, or a moving object with parts that change their spatial relations to one another or to other objects, such as a person or animal walking or running past other static or moving objects, a task addressed in [Hogg\(1983\)](#).

So the forms of compositionality involved in perceiving static structured scenes are much simpler than the forms required for perceiving moving objects, especially moving objects composed of moving objects.

Note on change-blindness

[Change blindness](#) is a well known phenomenon that occurs when humans fail to notice changes occurring in scenes or images of various sorts. It is widely thought that change blindness is puzzling, and needs to be explained. This is back-to-front! What many fail to realise is that **change detection** is what needs to be explained. Since this requires comparing a current state to one or more past states of the environment, it requires comparing something that exists with something that no longer exists. This can't be achieved simply by storing records of sensory stimulation (e.g. retinal stimulation), and then checking for changes between current stimulation and recent records of stimulation, because many changes in retinal stimulation are caused by eye motion, e.g. saccades, or visual tracking of a moving object. In that case there is a problem of explaining how absence of change in the scene is detected. (Left as an exercise for the reader.) This illustrates the fact that compositionality can occur in processes as well as in static structures. Mobile organisms, static organisms perceiving changing scenes and mobile organisms perceiving changing scenes, all need information process-mechanisms able to deal with different forms of compositionality, including composition of motion in the environment, motion of the observer and motion of fixation point. The achievements of biological evolution in addressing these challenges are spectacularly successful (e.g. in birds flying through trees), but barely understood so far, and not yet replicated in robots as far as I know. Is there a drone that can fly between branches and land on one, in a tree on which it has not been trained?

Compare: https://en.wikipedia.org/wiki/Change_blindness

Alternative forms of compositionality

In more complex cases, structures need to be represented by graphs rather than trees, e.g. if some of the constituent trees share sub-structures, e.g. in some representations of two verb phrases with the same subject. (I am not claiming that brains of English speakers explicitly construct structures like these when they understand a story, a sentence, a clause, or a phrase. Such issues are beyond the scope of this paper. Some compilers and interpreters for computer programming languages construct intermediate trees and graphs, but not all.)

A biological example might be the fact that different genetically influenced mechanisms of navigation in different environments can be used in organisms sharing some abstract design

features and needs (e.g. requiring various resources at different locations at different times) and differing in their concrete physical features and abilities, e.g. using different sensors and motors, or depending on land based propulsion, water-based propulsion or flying, for their seasonal migrations or other journeys.

Just as the semantic content of a complex sentence may be built up from the semantic contents of various parts, parts of parts, etc. so a complex navigational competence may be provided by composition of parts or features supplying different sub-competences that together contribute to the competence of the whole organism (or subsystem of the organism, such as the navigational subsystem).

That could in some cases result from a part of the genome being somehow derived by abstraction from special cases that evolved earlier. The derived abstraction would include "gaps" that might have different fillers in different individuals, or in different species.

The need for such abstraction could, for example, arise in species that change their size and shape during development. In that case the control subsystems using current goals and sensory information to decide on motor signals would have to be replaced continually, unless the control mechanisms can be divided into an abstract fixed design for a mechanism that is combined with changing parameters that vary with the size and shape of the organism as it grows.

The need for that type of control abstraction during development is evident in all organisms that change their size, shape, weight, weight distribution, strength, speed of motion and also their needs, e.g. for different kinds of food.

An alternative to use of such parametrised designs employing compositionality, is a genome that specifies a succession of designs for an individual organism and then arranges for the organism to revert to a chemical soup prior to instantiation of the next design, much as some species have a larval stage that is followed by a pupal stage after which a later design emerges with the chemical components reused in new structures, for instance in butterflies and other insects. I suspect this alternative is biologically (chemically? physically? mathematically?) impossible for organisms that (approximately) maintain their overall structure but undergo considerable (continuous) changes in overall size and ratios between sizes of parts, e.g. reptiles, mammals, birds, fish, and others.


This suggests that for such organisms evolution discovered the deep advantages of abstraction in part of a genome that can specify a structure whose details can change during development, in something like the way a human engineered control system can include control parameters that vary according to circumstances, e.g. allowing the rate of acceleration or deceleration of an automated vehicle to vary with the load or with sensed friction parameters between wheels and road or rails.

The next section indicates how reliance on a succession of designs for complete organisms during development of individuals can be avoided in a "meta-configured" genome, which has genetic specifications at different levels of abstraction that are instantiated at different stages during development of individuals, using parameters acquired during earlier stages of development. This generalises the notion of a design with parameters: the parameters for part of a design may be complex structures built during earlier stages of development while interacting with the local environment.

Waddington's "Epigenetic landscape"

Epigenetics is the study of processes and mechanisms by which genes influence the development of individual members of a species. A well known metaphor for this process as partly like a ball rolling down a landscape providing various routes from the high level starting point to a some lower level terrain was depicted as follows in 1957 by the biologist, C.H. Waddington, labelled "The Epigenetic Landscape". His theories evolved over time. I am using this merely in order to contrast it with the idea of a meta-configured genome, summarised below.

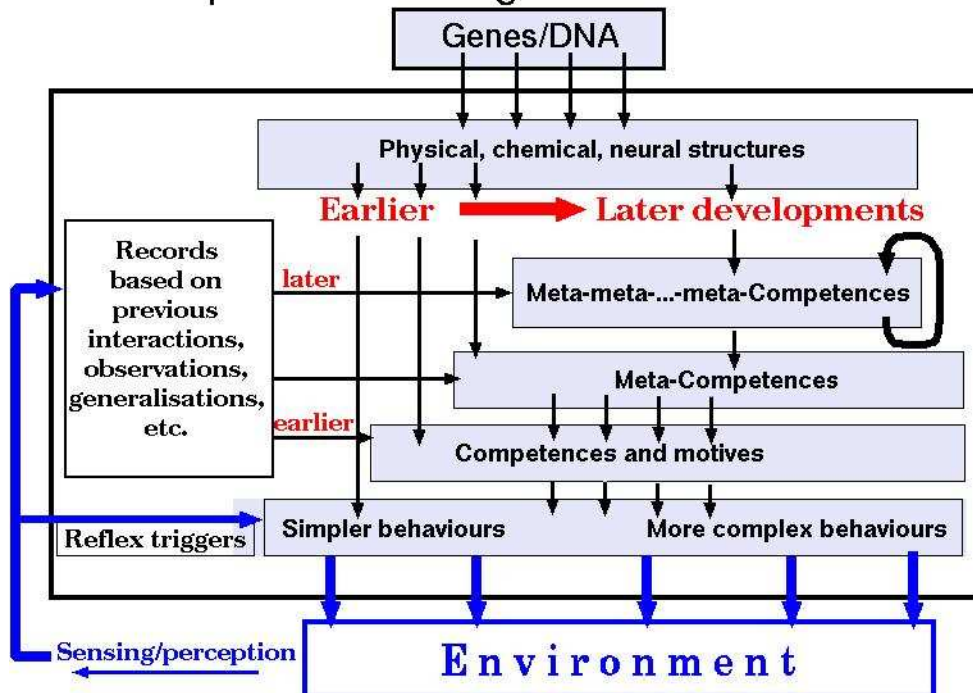
Figure Landscape



The main difference can be summed up as follows: on the "Meta-configured" landscape, as the ball rolls down the landscape what it finds enables it to change the structure of the rest of the landscape, modifying the options available thereafter (for that individual), partly on the basis of effects of its previous trajectory. Think of how the early years of learning a particular language alter language-development competences thereafter. Similarly the toys and games that a child, or other intelligent animal, learns to interact with or take part in may significantly alter the kinds of things that can be learnt, or the ease with which they can be learnt later. This is indicated crudely in [Figure Meta](#) below, though most of the details will have to be provided by future research.

A meta-configured genome’s use of compositionality

Figure Meta
Multiple routes from genome to behaviours



[Figure Meta](#) above, based on ideas in [Chappell & Sloman\(2007\)](#), shows (crudely) how staggered gene expression can allow relatively abstract/schematic results of later gene expression to be instantiated using information gained during earlier interactions with the environment -- not necessarily using any standard form of learning. The box on the left containing "Records" is a gross oversimplification -- earlier results of interaction with the environment will feed into changes of *design*, e.g. development of grammars, not merely historical records.

(A slightly clearer presentation of these ideas can be found in a 9.5minute video [Video 2] on the Meta-Configured Genome in this list

<https://www.youtube.com/playlist?list=PLYC-dSilAaYa6Mk1g6hBGUyqCwrlvyOWB>)

Earlier stages of gene expression are indicated by left-most black downward arrows, and later stages by downward arrows to the right. The earlier stages are mainly determined by the genome and control not only early growth patterns but also forms of instinctive behaviour e.g. consuming nutrients or avoiding harmful entities.

A well known phenomenon, imprinting, was discovered and studied by Konrad Lorenz, although it had been known and used by some farmers earlier. In some "precocial" species newly hatched individuals are able to move themselves, and to see things moving in the environment. It is important for them to follow the parent, who knows where to get food, etc., and in some cases will respond to their needs. Instead of the genome encoding precise details of the parent's features, it specifies a mechanism for getting the relevant details after hatching, from the first nearby large object seen to move, normally their mother in the case of geese, ducks, chickens, etc. This process is labelled "imprinting" (though the word also has other uses in biology).

Lorenz famously demonstrated that newly hatched geese would imprint on him and the tendency to follow him persisted across extended development, and in a wide variety of circumstances, as shown in this video

<https://www.youtube.com/watch?v=JGyfcBfSj4M>

This mechanism makes it unnecessary for the genome to include the detailed information about appearance of the parent, and avoids the (unsolvable?) problem of developing different genetic encodings for mechanisms for recognizing the right adult to follow, although it does require production of a mechanism capable of very rapid extraction and storage of the appropriate visual information details required to identify the parent, viewed in different contexts, and from different distances at different angles: no mean achievement.

During early stages, in addition to growth and feeding, [information](#) may be acquired, transformed, and then used during later stages of gene expression in combination with newly expressed genetic information.

Various stages in language development provide well-known examples, where later syntactic competences are in part controlled by information acquired during earlier stages, with simpler syntactic competences. Related ideas are in [Karmiloff Smith \(1992\)](#), discussed in

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

Depending on the complexity of the adult organism and the amount of growth and change required during development, several "staggered stages" of development may be required to produce the adult organism. For example, puberty normally includes both physical changes and motivational changes, whose precise form may depend not only on the genome but on information gained from the local culture.

In some species, even if physical development ceases after full adult size has been reached, information processing capabilities (e.g. perceptual skills, motor-control skills, learning capabilities, genetically triggered motive generators and motive comparators) may continue to develop, partly influenced by information acquired from the local culture. The effects are not always biologically useful: for example late developing meta-motivators related to discovering and adopting types of motivation that work well for more experienced conspecifics may be "hijacked" into fashion-industries that mostly benefit the fashion controllers.

Unfinished note:

In humans, sexual differentiation normally continues after adult size has been reached, including motivational mechanisms, partly because of the complexity and variety of processes produced in differing environments, that the genome seems to be able to cope with.

An example may turn out to be delayed expression of meta-competences required for contributing to education of younger members of the species, or contributing to management decisions for a social group, with motives relating to such decisions developing only after significant amounts of relevant information about the current situation have been acquired. A similar mechanism might produce late developing motives and skills in women past child-bearing age, so that their accumulated wisdom and experience feed into needs of larger groups than just the family.

(I have the impression that similar (simpler?) ideas were proposed in [Davies \(2012\)](#).)

Layered genome expression and use of construction kits

[Figure Meta](#) indicates how development under the combined influence of the genome and the environment involves cascaded, staggered, developmental trajectories, with later processes influenced by results of earlier processes in increasingly complex ways. Depending on the level of abstraction/generality encoded in later stages of gene expression, this kind of design can accommodate the need for member of the species to develop and thrive in a wide variety of environments.

This generality can be enhanced if the developing individual reaches new stages not merely by building on what has been acquired through interaction with the environment during early development, but also through guidance from, or imitation of older individuals (e.g. older conspecifics or parents).

As evolution progresses, for some species the mechanisms of evolution may become increasingly complex and varied, partly because of the use of increasingly sophisticated evolved construction kits during development.

In [Figure Meta](#), early genome-driven learning from the environment occurs in loops on the left. Downward arrows further right represent later gene-triggered processes during individual development building on and modulated by results of earlier learning as indicated by the horizontal arrows from left to right.

For example, whereas earlier processes might lead to learnt useful syntactic patterns available for linguistic communication, the later processes can use the products of earlier learning to create more sophisticated grammatical and semantic mechanisms supporting a richer variety of constructs -- illustrated by the child who once said to me, in defence of a proposal to go on a picnic in mid-winter "But today might be much more hotter than it usually be's", clearly expressing a complex thought with a newly created syntactic structure, using previously created linguistic competences to produce a type of sentence structure that obviously went beyond anything he had heard other speakers say.

In human language development, such creative compositionality at the semantic level is later modulated by creation of additional mechanisms for exception handling and for fine-tuning syntactic relationships. (This is sometimes referred to as "U-shaped" language learning: linguistic accuracy and syntactic complexity both develop for a while then there's a spurt in power accompanied by reduced accuracy, followed by later developments that allow syntactic irregularities in the language to be handled correctly. These are examples of what Karmiloff-Smith called "representational redescription".)

Non-quantitative semantic competences

As the example of language learning (among many others) shows, it would be a mistake to assume that all information is expressed in numerical measures, collections of measures, or functional relationships between measures. Much of the information acquired, created and used is structural, e.g. using syntactic structures in the case of linguistic competences, or plan construction, and topological or geometrical structures in visual competences. It is very likely that brains of humans and other intelligent animals create a much wider variety of internal structures with compositional semantics than scientists have so far discovered.

This is very different from designs now fashionable in some AI circles, using multiple layers where the basic principles of learning (using statistical evidence) are the same at all levels. For example, the later uses of linguistic competence in children not only to communicate but also to think, plan, reflect, and in some cases write stories or develop theories, cannot be based on the same learning mechanisms as the very early forms of learning used to produce sounds or signs and to recognise sounds or signs produced by others.

Similarly, it can be argued that later forms of perceptual competence, e.g. visual competence make use of increasingly complex non-numerical (e.g. topological) structures and relationships, that overlap with the linguistic semantic competences, and can be used in perception of complex relationships such as supporting, obstructing, moving out of the way, combining to form a larger structure, etc. This requires individuals to have not only percepts and action control structures but also increasingly complex plans and designs for novel objects or actions, built out of simpler versions either previously learnt or constructed when needed.

Evolution of ancient mathematical abilities

I suspect that the mathematical discoveries of ancient geniuses such as Archimedes, Euclid, Zeno, and many others before them depended on the use of newly created forms of perception, representation, goal formation, plan formation, and reasoning about impossibilities and necessary consequences, e.g. the impossibility of pushing a stick through a gap between two vertical surfaces if the stick is held horizontally and parallel to the surfaces.

If such an experienced impossibility is encountered, compositional creativity could make possible discovery of a solution to the problem by holding the stick vertically, or holding it horizontally but pointing into the gap between the surfaces rather than parallel to the surfaces. See [Sloman\(2007-2017\)](#) for examples of the importance of perceived changing geometrical and topological relationships when interacting with the environment.

The ability to represent, detect, and make use of cases of *impossibility* and *necessary consequence* can vastly reduce the need to use trial and error or empirical testing in achieving practical goals. I suspect that various restricted forms of these abilities are essential for various kinds of intelligence in non-human animals (e.g. squirrels) who are unable to communicate what they are doing or why it works, but nevertheless can acquire the key information and use it, though at present we still know very little about the details of these processes and the forms of representation and learning used.

Standard "neural net" brain mechanisms manipulating probabilities derived from statistical evidence cannot make discoveries regarding *necessity* and *impossibility*: characteristic features of mathematical discovery, as pointed out long ago in [Kant\(1781\)](#).

I suspect still unknown subsets of those competences are shared with other intelligent animals, in reasoning about spatial affordances and how they can change as a result of actions by the thinker, or another animal. A messy, and growing, collection of examples involving recognition of possibility, impossibility and necessity can be found in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>.

An example class of abstractions involving compositionality not discussed here, is the collection of navigational needs and abilities that are common across different forms of perception and motion, using different sensors and effectors.

I suspect that without such sharing of abstractions, supported by specially evolved "construction kits" [Sloman\(2017\)](#), evolution could not, in the time and space available, have produced so many different relatively successful species on this planet. (About a trillion species??)

The ideas are still being extended here

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

The UKCRC Grand Challenge on "Unconventional Computing" is also relevant:

<https://www.cs.york.ac.uk/nature/gc7/>.

The evolutionary significance of all this includes the possibility of dividing a complex genome into specifications of various abstract features that allow many different instantiations and specifications of particular combinations of features that instantiate those abstractions in highly context-sensitive ways. The compositionality of the various competences is crucial to the flexibility and creativity of

some organisms, including humans. Without this, it would not be possible for the human genome to support creation of several thousand widely varying languages for thinking and for communication, using very different physical media although no human seems to be able to become expert in more than a tiny subset.

However, the creativity of evolution itself underlies all the other examples, in ways that are not well understood.

Compositionality in natural language

In 1971 Winograd's PhD thesis (later published as (1972)) rapidly achieved international fame with its demonstration of the power of "heterarchical", as opposed to hierarchical, mechanisms for interpreting complex sentences, and the need to make use of information about syntax, semantics and facts about the world when understanding natural language communications, e.g. working out the difference between the referents of "they" in these two sentences:

The city councilmen refused the demonstrators a permit because they advocated violence.

The city councilmen refused the demonstrators a permit because they feared violence.

Recently this idea has been used as the basis of the so-called "Winograd schema challenge" to test comprehension abilities of natural language systems. Each schema includes two assertions and a potentially ambiguous pronoun, such as "they" or "it", as in the example above. The challenge is to specify the referent of the pronoun in each sentence. This requires not only understanding the compositional semantics of English but also using factual information about who is likely to advocate or fear violence. For more details see:

<https://cs.nyu.edu/faculty/davise/papers/WinogradSchemas/WS.html>

Perhaps a more dramatic illustration of the use of multiple sources of knowledge in dealing with linguistic compositionality is the ability to determine how the referent of "who" (i.e. the answer to the question) in these sentences changes as the sentence is extended:

"Who is the author of the book, *The Iron Lady: a biography of Margaret Thatcher*"

"Who is the author of the book, *The Iron Lady: a biography of Margaret Thatcher*, most likely to have been taught by?"

"Who is the author of the book, *The Iron Lady: a biography of Margaret Thatcher* most likely to have been taught by before leaving primary school?"

"Who is the author of the book, *The Iron Lady: a biography of Margaret Thatcher*, most likely to have been taught by before leaving primary school, in the town where she was born?"

and using a different branch, into the future:

"Who is the author of the book, *The Iron Lady: a biography of Margaret Thatcher*, most likely to cause to lose his seat in parliament?"

Answering these questions requires successively larger amounts of knowledge about the life of and related facts about the author, as well as general facts about the culture, the educational system, the history of a particular school and geographical relationships, and in the last example the current

situation in various parliamentary constituencies, not mentioned explicitly.

Determining who is referred to requires successively more complex logical and mathematical reasoning about the syntactic structure, the semantic content of the sentence and relevant factual information, reasoning that is normally done unconsciously at high speed by a human reader, using combinations of constraining information, i.e. information about necessary connections and impossibilities, that cannot be derived from statistical evidence.

The longer such a sentence is, the less likely it is to be structurally and semantically similar to a sentence in a training corpus, and the more likely it is that the disambiguation will require reasoning based on previously acquired facts that are rapidly identified as relevant to the question even though they have never previously been used in this context. This is not reasoning based on learnt statistical regularities such as might be used in a "deep learning" AI system, but reasoning closer in kind to ancient mathematical reasoning about spatial structures, including reasoning about impossibilities and necessary connections rather than probabilities derived from statistical observations.

"Mostly likely" in normal human discourse does not refer to numerical computations based on probabilities derived from statistical data. It is a much more subtle concept referring to a *partial* ordering of relative likelihoods normally derived from structural information about influences and constraints, not observed frequencies. Justifying that claim would require a much more complex discussion combining topics on which my own thinking is still incomplete. If two trees are similar in structure but tree A is much taller than tree B, can you answer the question: If Joe attempts to climb to the top of tree A is he more or less likely to damage himself than if he attempts to climb to the top of tree B? Why don't you need to collect statistics regarding people climbing and falling off trees in order to answer this?

Intensional and extensional compositionality

Compositionality in sense (connotation, Sinn) and reference (denotation, Bedeutung)

***Sinn vs Bedeutung*, in Frege's terminology.**

A complication that as far as I know has not been noticed by AI researchers, but is familiar to (some) philosophers is that not every syntactically well-formed expression in a language has the kind of semantic content that determines a denotation, or reference. In some cases that is because something assumed to exist doesn't actually exist as in "The cafe in which Einstein and Euclid first discussed geometry". That expression has a clear sense, and if Einstein and Euclid had overlapped in time and space and discussed geometry in a cafe somewhere the expression would also denote something, or refer to something.

In other cases reference/denotation may fail because of the infinite recursion involved in identifying the denotation, e.g. in "The father of the subject of this sentence is a teacher" [Sloman, 1971a](#).

Other cases of failure of reference are familiar in ancient philosophical problems, such as how to determine whether everything in the universe is moving in a certain direction, including all measuring devices, or working out the answer to Wittgenstein's question "What time is it on the sun?" or related cases with different causes of denotational failure e.g. "What time is it at the centre of the earth? or a billion kilometers above New York?" There are also mathematical examples, such as "The largest prime number". A collection of pictorial examples involving various kinds of

impossibility and necessity is in [Sloman\(impossible\)](#).

The ability to create and understand such collections of sentences requires what could be called meta-compositional competences: a kind of meta-knowledge about one's own compositional competences. I suspect that is also a key feature of abilities to make discoveries in mathematics, including research that leads to new sub-fields of mathematics, for example as Group Theory grew out of reflection on aspects of arithmetic, among other things.

It may be that that sort of capability is itself a special case of a type of creative ability that evolved much earlier and allows compositional expertise to be gained in exploring several domains, e.g. domains of structures that could be built with particular materials, e.g. types of clothing that could be made from animal skins, types of useful structure that could be made by assembling stones, boulders and branches, and types of action that could be generated by combining physiology with structures in the environment, e.g. climbing trees, throwing stones, throwing stones at particular targets, making throwable linear objects from straight portions of trees, etc.

Reasoning about appearance changes caused by motion

Another, much more subtle requirement for compositionality, which could have been a factor in development of geometrical and topological reasoning abilities prior to the discoveries made by great ancient mathematicians, is the ability to reason about how appearances of seen objects must change as either the perceiver's viewpoint changes or the objects themselves move either relative to other objects or relative to the perceiver, or relative to perceivable parts of the perceiver (e.g. hands, tentacles, an elephant's trunk, jaws, among many others). This includes reasoning about whether object parts or distinct objects will be separated, or brought together, resulting in perceived changes of shape and changes of relationships between shapes.

Many intelligent species, including nest-building birds and hunting animals need to be able to detect and reason about such changes in perceived appearance caused by motion of the perceiver or one or more parts of the perceiver, or motion of one or more perceived objects.

A special subset of such abilities involves the ability to work out that a particular appearance or appearance change is *impossible* in certain circumstances, or will *necessarily* occur in other circumstances. Examples are the changes in visibility of portions of a room seen through an open doorway as you move closer to or further from the doorway, or move left or right while looking through the doorway. More examples are discussed in [Sloman\(2007-2017\)](#) and [Sloman\(impossible\)](#).

It seems that such creative uses of compositionality are found not only in humans but also in a subset of non-human species, e.g. nest-building birds, and also in development of new behaviours, such as developing hunting skills in mammals, reptiles, birds, etc. A spectacular example in humans is the ability of deaf children to create their own new sign language, demonstrated in this video <https://www.youtube.com/watch?v=pjtioIFuNf8>. See also [Senghas\(2005\)](#).

The mechanisms used by the deaf children to create their language were not created by them, but by evolutionary processes (mostly still unknown, as far I can tell).

Spider creativity

In other cases, the creativity is not achieved by individuals but by evolutionary mechanisms, e.g. the mechanisms that produced web-creating competences in spiders, that seem to be mostly innate [Vollrath et al.](#), unlike the creative route-planning of the Portia spider in creating a novel route to a location above its prey [Tarsitano and Jackson \(1997\)](#) [Tarsitano \(2006\)](#).

The ability to rule out combinations of possibilities also plays a role in perception of and reasoning about physical structures and processes in the environment, using geometrical or topological impossibilities and necessary connections, e.g. transitivity and asymmetry of many relations expressed by prepositions, such as "inside", "above", "further than", etc.

Such reasoning abilities, originally developed for reasoning about affordances, in some cases based on aspects of 2D projections of 3D relationships, may have been important precursors of the amazing discoveries reported in Euclid's *Elements*.

Compositionality in biological niches

Added 4 Oct 2018

A biological niche is not a geographical location, but something much more abstract, a set of requirements for a functioning design, as explained in [Sloman\(1995\)](#) and other papers to be added here.

This section will be expanded to include pointers to work on evolutionary trajectories in niche-space and design-space. During evolution, sets of requirements for members of a species keep changing, partly as a result of changes in various more or less abstract features of the environment, including changing facts about competitors, prey, predators, and various resources including water, various foods, shelter, etc., and partly as a result of changes in the individual, e.g. growth in size, in skill, in knowledge, or injury, illness, etc.

Some of the individual changes result from evolutionary changes, others from individual development, and others from learning and cultural changes.

During those changes the designs instantiated by individuals can also change, as a result of evolution, growth, development, learning, and in some cases injury or genetic abnormality.

All this implies that there are complex and tangled relationships between changing biological designs (including physical changes and changes in cognitive competences) and changing niches (sets of requirements). Composition of all these changes can produce fluctuation in population sizes as well as changes in habitat and behaviours. In the case of humans, changes in knowledge, available technology and motivation can interact to produce dramatic, large scale effects, such as the effects of the internet and mobile electronic devices.

All this is further complicated by the ability of humans, unlike any other species on this planet, to be brainwashed by religions, including dreadful ethical systems leading to death and destruction on large scales. I don't know whether scientific understanding of the forms of compositionality involved can lead to new useful forms of control.

To be expanded

Cross-species/Symbiotic compositionality

19 Jul 2019: To be expanded

Many forms of biological compositionality involve two or more species -- e.g. humans and the multiple microscopic organisms without which they could not survive and many forms of symbiosis. Also forests, etc. etc.
... To be expanded...

[BACK TO CONTENTS](#)

CONCLUSION

This paper presents more questions than answers, though sometimes finding the right questions is a necessary precursor to a stream of good answers. I hope to have formulated at least some of the right questions, with the help of many colleagues and students, over half a century, and help from various great philosophers and scientists from beyond the grave (including Kant, Darwin, Frege, [Schrödinger\(1944\)](#), [Craig\(1943\)](#), [Turing\(1952\)](#), and others).

The right questions include questions relating to the nature and origins of biological spatial reasoning abilities, and their relationships to the mathematical reasoning abilities of ancient mathematicians who made important discoveries about geometry and topology. These are related to questions about spatial reasoning in non-human animals, the development of spatial reasoning in human infants and children, and the mechanisms that future robots will need in order to be capable of the kinds of spatial intelligence found in humans and other animals. That still seems far off, and may require new forms of computing machinery. It isn't clear to me that digital computers provide adequate forms of information processing. Statistics-based neural mechanisms are also inadequate insofar as they are incapable of discovering, or even representing the kinds of necessity and impossibility that characterise mathematical discoveries, as Immanuel Kant noticed, two and a half centuries ago.

NOTE

The important questions are not questions about empirical regularities in the behaviour, or mental processes, in animals (human or non-human). Those are shallow questions whose answers are subject to many contingencies affecting research results, including choice of individuals to test (compare [Andreski\(1974\)](#)). The deeper questions are about what sorts of things are *possible* or possible for particular types of individual and *what makes those things possible* (explains the capabilities) of those individuals. This requires specifications of possible mechanisms with the required generative power, rather than discovery of regularities in behaviour. (This distinction has been understood by theoretical linguists since the mid 1960s, and by some AI researchers, e.g. Max Clowes. For more detail regarding explaining possibilities see [Chapter 2](#) of [Sloman\(1978\)](#).)

There are different questions about foundations of geometry discussed in [Corfield \(2017\)](#). At some future time it would be interesting to investigate in more detail the relationships between the two sets of questions.

This paper was written for a conference on compositionality with a strong emphasis on category theory

<http://events.cs.bham.ac.uk/syco/1/>

However, I have not found a role for category theory in my attempts to explain spatial cognition in humans and other animals, and the type of spatial cognition that made possible the discoveries made long ago about geometry and topology.

I am not sure how relevant category theory is to this project, but I am sure that the project needs help not only from biologists, but also from mathematicians and computer scientists with deep biological knowledge.

The ideas developed with biologist Jackie Chappell and presented in [Chappell & Sloman\(2007\)](#), spell out in a little more detail how epigenetic mechanisms and processes can be conceptualised as involving a mixture of top-down and bottom-up information flow during individual development as indicated schematically in [Figure Meta](#). But there is far more work to be done, some of it sketched in the M-M project. I have the impression that a great deal of ongoing experimental and theoretical research in biological departments has already gone beyond the ideas presented here, but in narrowly focused contexts, e.g. particular reproductive or developmental processes in particular types of organism.

Individual organisms can all be thought of as decomposable into structure plus components, but work in the M-M project shows that this is also true of many of the products and mechanisms of evolution, not just individual organisms.

Moreover, since it is a core feature of life that organisms can grow in size, change their shape in various ways, change the spatial relations of their parts, and change spatial relationships to other organisms and physical objects, a deep explanatory theory about how all that is possible, must say something about the space in which such structures and processes occur. That is not addressed by work on foundations of mathematics that search only for notational structures and processes, usually using discrete notations as in logic and algebra.

It may turn out that finding answers to the questions posed here will require, and may eventually provide, a new kind of foundation for a major subset of mathematics, not based on logic, algebra and set theory, or restricted to use of linear discrete notations, but rather based on mechanisms for reasoning about possible and impossible structural changes in various complex entities, some discrete and others continuously variable, and the necessary consequences of those changes, especially when they interact.

For instance changing the shape, size and location of a polygon can affect the number and locations of possible points common to the boundaries of a circle and the polygon.

Turing's 1952 paper on morphogenesis described mechanisms able to generate a subset of such changing spatial structures. Was he thinking about moving on to something more general? I've tentatively labelled the hypothetical target a "Super-Turing Multi-Membrane" machine, and will later add links here to some draft ideas.

That may mark a major break with "foundational" work in mathematics that uses only discrete formalisms and structures, even when discussing and reasoning about continuity.

Recent work by Dana Scott (2014) on geometry without points may be relevant. A video recording of a presentation in Edinburgh in 2014 is available here (with many continuous gestures and drawing actions, in addition to discrete spoken words and written symbols):

<https://www.youtube.com/watch?v=sDGnE8eja5o>

Even if physical matter is discrete (e.g. point-like) at the lowest level, perhaps motions and other changes of matter are not? For example, it may be that configurations of matter (e.g. molecules) can change chemical bonds only discretely (as discussed by [Schrödinger\(1944\)](#)), but molecular motions such as moving together or apart, folding twisting, etc. seem to require a continuous space in which such motions are possible.

If neither space nor matter is continuous then rotation of very large objects must produce anomalies at great distances from the centre of rotation.

This is a long term difficult project -- too difficult for me. Collaborators, or tutors, would be very welcome.

THANKS

To biologist Jackie Chappell, for joint work on the idea now labelled "The Meta-Configured Genome", first presented jointly at [IJCAI in 2005](#), then extended in our [2007 IJUC paper](#), and continued collaboration thereafter.

<https://www.birmingham.ac.uk/schools/biosciences/staff/profile.aspx?ReferenceId=5282&Name=dr-jackie-chappell>

To [Jamie Vicary](#), for organising the [SYCO-1](#) symposium, and encouraging me to submit a paper, which triggered a new shift in my thinking, and to the reviewers for their encouraging comments on the submitted version of this paper. To Achim Jung for encouragement, useful comments and suggestions. To a collection of students and colleagues over many years, for criticisms, suggestions and new ideas, including Geoffrey Hinton when we worked together at Sussex University in the 1970s, where he tried, and failed, to convince me that a trained neural network associating images with learnt labels could replicate or model perception of a ballet performance with constantly changing groups with changing patterns of relationships while seeing the moves of the leading dancer or pair of dancers, and Alan Mackworth who introduced me to the idea of a mind like a room full of loaded mouse traps and ping-pong balls, reacting to a single new ping-pong ball thrown in. I hope this paper is not too much like that.


APPENDIX A

A tutorial example illustrating linguistic compositionality (Early draft: to be revised)

I hope this will provide some useful background for readers unfamiliar with language processing concepts and techniques shared between Linguistics, AI, and Psychology of language, although there seems to be no body of theory that everyone agrees with.

Here is a crude picture of a scene containing blocks on a table and a (schematic) hand at the top:

Figure Blocks



A question is typed in about the contents of the scene:

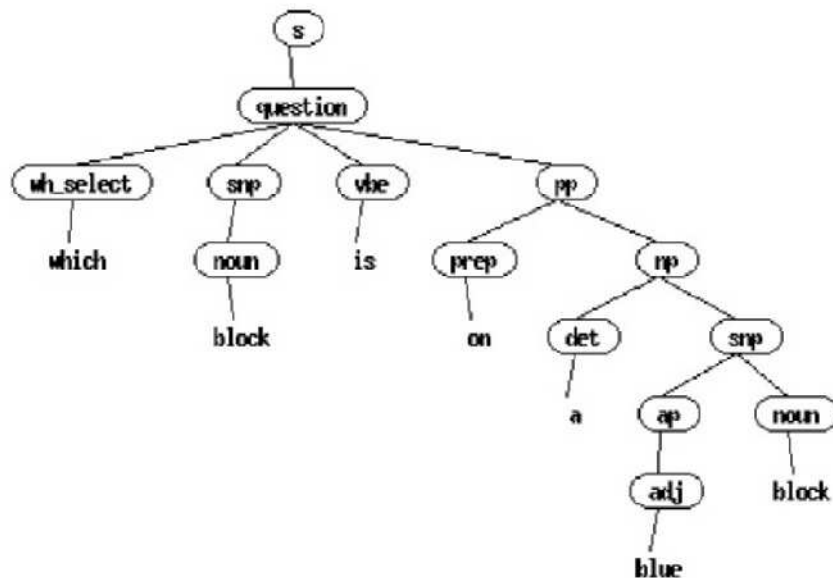
```
which block is on a blue block?
```

What mechanisms enable your brain to answer that question?

The program cannot see the picture, but instead uses an internal "logical" language in which a collection of facts about the world has been provided. It also has information about possible actions that can change the world, and it is able to represent such changes by changes in its list of facts.

The program has also been given a simple grammar plus a simple vocabulary, for a tiny subset of English, along with a sentence parser for a toy "demo" grammar parser for that subset, which uses that information to analyse a sentence typed in (in this case a question), and work out its grammatical structure represented here as a "tree structure":

Figure Parse tree



That tree is represented internally by the program as a list of lists:

```
** [s [question [wh_select which]
               [snp [noun block]]
               [vbe is]
               [pp [prep on] [np [det a] [snp [ap [adj blue]] [noun block]]]]]]]
```

The labels in the list distinguish various sub-components such as prepositions (prep), nouns

(noun), adjectives (adj), determiners (det, e.g. "a", "the", "an"), noun phrases (np), simple noun phrases (snp), adjective phrases (ap), prepositional phrases (prep), and so on. To the right of each left square bracket [is a label, indicating the syntactic category found, e.g. this is a sentence (s) which is a question (question) of type "wh_select" (the answer should select something), etc.

In this program, the structures have logical forms and a collection of procedures that allow the program to interrogate its internal "database" containing facts about the objects and their relationships, in order to obey commands or respond to questions, as in the reply below to the above question:

That block is the small red block.

In addition to answering questions this toy program can also obey commands expressed in simple English e.g.:

Put the block on a block on a red block on the small green block

And it can also detect that a command cannot be obeyed, for example

Put the small block on a blue block on the small red block

which fails when it tries to put the small red block on the small red block! A more intelligent program would discover the impossibility earlier, either by generalising from failed attempts, or, if it has appropriate forms of meta-cognition, it may discover (how?) the mathematical fact that "on" (immediately above and touching) is an antisymmetric, irreflexive relation. Therefore, achieving a state where something is on itself is impossible.

What sort of robot brain would be able to discover that if a command has semantic content in a particular context amounting to

Put X on X

then obeying the command is impossible, as opposed to merely discovering that its attempts fail? Statistical evidence cannot prove impossibility: "P is impossible" is quite different from "The probability of P is 0%".

A simpler program, or a less intelligent animal, could continue indefinitely trying and failing to perform an action, without having the meta-cognitive ability to detect that it **cannot** succeed. How do human brains make such discoveries? Compare a wasp repeatedly trying to escape through the closed portion of a window, while another portion of the same window is open. (There are many online videos of animals, including some very young children, attempting impossible actions, as well as videos of squirrels, crows and other animals working out how to transform an impossible task to something achievable.)

All of these cases involve abilities to analyse structures (perceived or typed in, or internally created) in terms of their parts, parts of parts (... etc.) and relationships between parts. And the interpretation of a larger structure, e.g. a complete sentence, or a visual image of a scene, is created by assembling interpretations of parts and using them to derive the interpretation of the whole sentence, image, or other structure. That process may require creative problem-solving, e.g. if some alternative candidates for part of a solution need to be consistent with other alternative candidates for other parts.

More information about this toy program can be found here,

<http://www.cs.bham.ac.uk/research/projects/poplog/examples/#shrdlu>

And two video demos of the program at work (without spoken commentaries):

<http://www.cs.bham.ac.uk/research/projects/cogaff/movies/poplog/gblocks.webm>

<http://www.cs.bham.ac.uk/research/projects/cogaff/movies/poplog/gblocks1.webm>

(The package used for these demos is included with the linux version of the Poplog system.)

NOTE

Neither Winograd's program, nor our toy version described above and used in the demos, can see, or make any sense of the picture of the scene. (In Winograd's program pictures of 3D scenes were generated for the benefit of human viewers, whereas our "toy" version used only much simpler 2D projections.) The program accesses only a collection of facts expressed as combinations of symbols in an internal database originally provided by the programmer, and when the program "performs" actions the actions simply alter the database and the corresponding pictures. (Those limitations do not undermine the educational role of such a system.)

In principle such a program could be combined with a physical robot manipulating real blocks on a real table, but that would require a very much bigger program including visual and motor control abilities and abilities to run different virtual machines in parallel, such as an arm control subsystem, a visual subsystem, and a meta-cognitive subsystem reasoning about what is happening and what can and cannot happen.

It is not clear what sorts of mechanism would enable an advanced version of such a machine to repeat the discoveries about geometry and topology made by our ancient predecessors, e.g. Archimedes.

Winograd's program was developed in 1971 when computers were physically very much bigger than current desktop PCs but had tiny memory capacities (e.g. a megabyte memory was considered huge, whereas now a wrist-watch could have a larger memory!). The program also had very slow processors by 21st Century standards (e.g. 1Mhz was very fast). Many critics of early symbolic AI have no idea of the practical obstacles the pioneers faced. (They often present alternative theories, e.g. theories of "embodied" or "enactive" or "situated" cognition, without having any idea what would be involved in designing, testing, and debugging such a machine.)

A modern robot with far more memory and very much more processing power, may be able to use cameras in its eyes and analyse the images recorded to produce an internal database of information about objects on a table and their relationships, and use that information in forming an intention, making a plan, and carrying out the plan by controlling motors, or answering questions. A hungry squirrel trying to get at nuts in a bird-feeder will need to do something of the same general sort but with far more complex details and at much greater speed. (Search the internet for videos of squirrels and bird-feeders.)

Most electronic cameras currently record visual information in rectangular arrays of numbers, or in some cases concentric circular bands of numbers. Cells in animal visual systems do not have such regularities, so the tasks they perform are different.

The Kinect visual device does something completely different, using a "stereoscopic" visual system locating surface points in 3D space. I believe that for many purposes that is a less useful design for an intelligent robot visual system than one that uses projection into a 2D surface, e.g. because computations derived from changing 2D projections will often use less computing power to produce information that is very useful for taking practical decisions (e.g. whether the surface of object A obscures more or less of object B as you move your viewpoint to the right). This is one among very many examples of process compositionality that are useful for intelligent agents interacting with a complex physical world.

I suspect these aspects of process compositionality in perceptual and reasoning system in our ancient ancestors were core precursors of the cognitive mechanisms that allowed the great ancient mathematicians and engineers to produce their amazing mathematical discoveries and engineering constructions long before the development of coordinate geometry and designs for logical reasoning mechanisms.

How animal brains represent desires, intentions and environmental facts, and how those factual representations are related to visual mechanisms, planning capabilities, plan execution abilities, and in the case of humans linguistic abilities used when describing or understanding descriptions of spatial configurations, are still unexplained, as far as I know. (Try asking a few penetrating questions next time you see a robot demonstration.)

Yet it is clear from the observed capabilities of many animals that their brains are able to take in and make use of information about external structures and are able to create intentions and in some cases multi-step plans involving sequences of [possible](#) actions. (Compare the Portia spider mentioned [above](#).) For reasons not explained here, I do not think robots based on current AI techniques will match the spatial intelligence of squirrels, nest-building birds, orangutans, hunting mammals, or pre-verbal humans or ancient mathematicians, in the foreseeable future. This type of intelligence builds on competences based on internal structure manipulation, found in many intelligent non-human animals, as well as to humans in various physical environments, in various cultures, with various kinds and amounts of knowledge and experience. Some of our very distant ancestors must have had genomes with the [potential](#) to develop expertise in Euclidean geometry long before humans ever thought of perfectly thin, perfectly straight, perfectly parallel lines. This is totally different from currently fashionable statistics-based forms of intelligence.

I suspect our better informed ancestors will look back and regard deep learning theories and related probability-based models of intelligence, as comparable to the spread of mental metaphors related to steam engines and hydraulic power in the age of steam power. If you have only a hammer, everything is a nail. I hope that comment doesn't cause the reader to explode.

[BACK TO CONTENTS](#)

REFERENCES

Stanislav Andreski, 1974, *Social Sciences as Sorcery*
Jan, 1974 Penguin Books

Xerxes Arsiwalla, Martha Lewis, Mehrnoosh Sadrzadeh and Paul Verschure. (2018), Conscious Perception as a Non-Classical Language of Thought, in *Proceedings CAPNS 2018*
<https://sites.google.com/view/capns2018/accepted-papers>

Dana Ballard and Chris Brown, 1982, *Computer Vision*, Prentice Hall, Inc., Englewood Cliffs, New Jersey 07632, USA,
<http://homepages.inf.ed.ac.uk/rbf/BOOKS/BANDB/bandb.htm>

Jacob Beck, 2018, Do nonhuman animals have a language of thought?, in *The Routledge handbook of philosophy of animal minds*, Eds. K. Andrews and J. Beck, pp. 46--55, Routledge handbooks in philosophy, New York, NY, Routledge/Taylor & Francis Group,
<http://psycnet.apa.org/record/2017-44151-004>

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

Jackie Chappell and Aaron Sloman, (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>

Clowes, M. (1971). On seeing things. *Artificial Intelligence*, 2(1), 79-116.
[http://dx.doi.org/10.1016/0004-3702\(71\)90005-1](http://dx.doi.org/10.1016/0004-3702(71)90005-1)

Clowes, M. (1973). Man the creative machine: A perspective from Artificial Intelligence research. In J. Benthall (Ed.), *The Limits of Human Nature*. London: Allen Lane.

Coates J, Umm-E-Aiman, Charrier B (2014) Understanding "green" multicellularity: do seaweeds hold the key? *Frontiers in Plant Science*, Dec 2014
<http://doi.org/10.3389/fpls.2014.00737>

S. B. Cooper & J. van Leeuwen (Eds.), *Alan Turing - His Work and Impact* (p. 849-856)
Amsterdam: Elsevier.
(Table of contents available here
<http://www.cs.bham.ac.uk/~axs/amtbook>)

David Corfield (2017) Reviving the philosophy of geometry, in *Categories for the Working Philosopher* Ed. Elaine Landry, OUP <http://philsci-archive.pitt.edu/11809/>

Kenneth Craik, (1943), *The Nature of Explanation*, 1943, CUP, London, New York.

Davies, P. C. W. (2012). The epigenome and top-down causation. *Interface focus*, 2 1, 42-8.
<https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3262298/> (Online 14 September 2011)

J.A. Fodor, 1975, *The Language of Thought*, Harvard University Press, Cambridge,

Frege, G. (1960). In P. Geach & M. Black (Eds.), *Translations from the philosophical writings*. Oxford: Blackwell.

Gottlob Frege, Sense and Reference, translated in *The Philosophical Review*, Vol. 57, No. 3 (May, 1948), pp. 209-230

Tibor Ganti, 2003. *The Principles of Life*, (Translation of the 1971 Hungarian edition), Eds. E. Szathmáry, & J. Griesemer, OUP, New York.

See the very useful summary/review of this book by Gert Korthof:

<http://wasdarwinwrong.com/korthof66.htm>

H. Gelernter, 1964, Realization of a geometry-theorem proving machine, in *Computers and Thought*, Eds. Feigenbaum, Edward A. and Feldman, Julian, pp. 134-152, McGraw-Hill, New York, Re-published 1995 (ISBN 0-262-56092-5),

<http://dl.acm.org/citation.cfm?id=216408.216418>

Also at

<https://pdfs.semanticscholar.org/2edc/8083073837564306943aab77d6dcc19d0cdc.pdf>

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

Seth G. N. Grant, 2018, Synapse molecular complexity and the plasticity behaviour problem, *Brain and Neuroscience Advances* 2, pp. 1--7,

<https://doi.org/10.1177/2398212818810685>

J.A. Hampton (2017), Compositionality and Concepts. In: Hampton J., Winter Y. (eds) *Compositionality and Concepts in Linguistics and Psychology. Language, Cognition, and Mind*, vol 3. Springer.

https://doi.org/10.1007/978-3-319-45977-6_4

Hodges, W. (2014). Tarski's truth definitions. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2014 ed.). Metaphysics Research Lab, Stanford University.

<https://plato.stanford.edu/archives/fall2014/entries/tarski-truth/>

David Hogg, 1983, Model-based vision: A Program to see a walking person, in *Image and Vision Computing*, 1, 1, pp. 5--20,

Huffman, D. A. (1971). Impossible objects as nonsense sentences. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence*, (Vol. 6, pp. 295-323). 22 George Square, Edinburgh, Great Britain: Edinburgh University Press.

Eva Jablonka and Marion J. Lamb, 2005, *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*, MIT Press, Cambridge MA,

S. Kaneff (Ed), 1970, *Picture language machines*, Academic Press, New York.

Kant, I. (1781). *Critique of pure reason*, London: Macmillan. (Translated (1929) by Norman Kemp Smith)

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press. (Discussed in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html>)

Catriona M. Kennedy, 2003, *Distributed Reflective Architectures for Anomaly Detection and Autonomous Recovery* PhD thesis, University of Birmingham

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

Kirschner, M., & Gerhart, J. (2005). *The Plausibility of Life: Resolving Darwin's Dilemma*. Princeton: Yale University Press.

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

Seth Lloyd, 2012, "Seth Lloyd on quantum life (Video talk)", Perimeter Institute for Theoretical Physics, 9 Oct 2012 <http://www.kurzweilai.net/seth-lloyd-on-quantum-life>

John McCarthy, 1977, Epistemological Problems of Artificial Intelligence, *Invited talk IJCAI 1977* <http://jmc.stanford.edu/articles/epistemological.html>

David Mumford(Blog), Grammar isn't merely part of language, Oct, 2016, Online Blog, <http://www.dam.Brown.edu/people/mumford/blog/2016/grammar.html>

Francis J. Pelletier, 2001, Did Frege Believe Frege's Principle?, *Journal of Logic Language and Information*, pp. 87--114, Vol 10, No 1, <https://doi.org/10.1023/A:1026594023292>

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

Science History Institute, 2017, on James Watson, Francis Crick, Maurice Wilkins, and Rosalind Franklin "These four scientists--Watson, Crick, Franklin, and Wilkins--co-discovered the double-helix structure of DNA, which formed the basis for modern biotechnology." Dec 2017 <https://www.sciencehistory.org/historical-profile/james-watson-francis-crick-maurice-wilkins-and-rosalind-franklin>

Dana S. Scott, 2014, Geometry without points, (A preliminary report on on-going joint work with Tamar Lando, Philosophy, Columbia University.) <https://www.logic.at/latd2014/2014%20Vienna%20Scott.pdf>

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>

[Some of my own publications on this topic are [below](#)]

Szabs, Z. G. (2017). Compositionality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2017 ed.). Stanford University. <https://plato.stanford.edu/archives/sum2017/entries/compositionality/>

Tarsitano, M. S., & Jackson, R. R. (1997). Araneophagic jumping spiders discriminate between detour routes that do and do not lead to prey. *Animal Behaviour* 53, 257-266.

Tarsitano, M. (2006, December). Route selection by a jumping spider (*Portia labiata*) during the locomotory phase of a detour. *Animal Behaviour*, 72, Issue 6, 1437-1442. <http://dx.doi.org/10.1016/j.anbehav.2006.05.007>

Simon W. Townsend, Sabrina Engesser, Sabine Stoll, Klaus Zuberbuhler, and Bickel Balthasar, Compositionality in animals and humans, August 2018, in *PLOS Biology*, Public Library of Science, Vol 16, 8, pp. 1-7, <https://doi.org/10.1371/journal.pbio.2006425>

Trettenbrein, Patrick C., 2016, The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift?, *Frontiers in Systems Neuroscience*, Vol 88, <http://doi.org/10.3389/fnsys.2016.00088>

A. M. Turing, (1950) Computing machinery and intelligence, *Mind*, 59, pp. 433--460, 1950, (reprinted in many collections, e.g. E.A. Feigenbaum and J. Feldman (eds) *Computers and Thought* McGraw-Hill, New York, 1963, 11--35), WARNING: some of the online and published copies of this paper have errors, including claiming that computers will have 109 rather than 10^9 bits of memory.

A. M. Turing, 1952, The Chemical Basis Of Morphogenesis, *Phil. Trans. R. Soc. London B* 237, 237, pp. 37--72, (Reprinted in *Alan Turing, His Life and Work* Elsevier, 2013).

Alan Turing's 1938 thoughts on mathematical reasoning (intuition vs ingenuity)
Discussed by Aaron Sloman, 2018
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html>
(also pdf).

Vollrath, F., & Selden, P. (2007, Sep 5). The Role of Behavior in the Evolution of Spiders, Silks, and Webs. *Ann. Rev. of Ecology, Evolution, and Systematics*, 38, 819-846. <http://dx.doi.org/10.1146/annurev.ecolsys.37.091305.110221>

Alastair Wilson, 2017, Metaphysical Causation, *Nous*
<https://doi.org/10.1111/nous.12190>

Winograd, T. (1972). Procedures as a Representation for Data in a Computer Program for Understanding Natural Language. *Cognitive Psychology*, 3(1).

Some of my publications related to this topic

Sloman, A. (1962). *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth* (DPhil Thesis), Oxford University. (Transcribed version online.)
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962>

A. Sloman, (1965), Functions and Rogators, in *Formal Systems and Recursive Functions: Proceedings of the Eighth Logic Colloquium* Oxford, July 1963, Eds. J. N. Crossley and M. A. E. Dummett, North-Holland Publishing Co, Amsterdam, pp. 156--175, <http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#rog>

A. Sloman, 1971, Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence, *Proc 2nd IJCAI*, London, pp. 209--226,
<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#1971-02> William Kaufmann,
Reprinted in *Artificial Intelligence*, vol 2, 3-4, pp 209-225, 1971,

Sloman, A. (1971a). Tarski, Frege and the Liar Paradox. *Philosophy*, 46(176), 133-147.
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1971-03>

Aaron Sloman, 1978

The Computer Revolution in Philosophy: Philosophy science and models of mind,
Harvester Press and Humanities Press, 1978. Revised edition (2001--) online:
<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/>

Aaron Sloman (1983), Image interpretation: The way ahead? In: *Physical and Biological Processing of Images* (Proceedings of an international symposium organised by The Rank Prize Funds, London, 1982.), Eds. Braddick O, Sleigh A, Springer-Verlag, Berlin, pp 380-401,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#57>

A. Sloman, 1984, The structure of the space of possible minds, in *The Mind and the Machine: philosophical aspects of Artificial Intelligence*, Ed. S. Torrance, Ellis Horwood, Chichester,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#49a>

A. Sloman, 1984-Clowes, Experiencing Computation: A Tribute to Max Clowes, in *New horizons in educational computing*, Ellis Horwood Series in AI, Ed. Masoud Yazdani, pp. 207--219,
(Online version with expanded obituary and biography.)
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61>

Aaron Sloman, What Are The Purposes Of Vision?, Presented at the Fyssen Foundation Vision Workshop, Versailles France, March 1986, Organiser: M. Imbert,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#58>

Aaron Sloman, 1995, Exploring design space and niche space, (Keynote), *Proceedings 5th Scandinavian Conference on AI*, Trondheim, IOS Press, Amsterdam,
<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#41>

Aaron Sloman, colleagues, and students October 1994-2015:

TEACH OOP: An introduction to Object Oriented Programming in Poplog

On OOP: <http://www.cs.bham.ac.uk/research/projects/poplog/teach/oop.html>

On Poplog: <http://www.cs.bham.ac.uk/research/projects/poplog/freepoplog.html>

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

Introductory overview of the Birmingham Cognition and Affect Project

by Aaron Sloman, colleagues, and students.

Aaron Sloman and Jackie Chappell, The Altricial-Precocial Spectrum for Robots,
in *Proceedings IJCAI'05*, Edinburgh 2005, pp. 1187--1192,
<http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>

Note: sequels include this invited paper with Jackie Chappell [IJUC 2007](#)

and work in progress: The Meta-Configured Genome (also mentioned [above](#)), and

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

Aaron Sloman, (2006--), Requirements for a Fully Deliberative Architecture, (technical report), <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/fully-deliberative.html>

Sloman, A. (2007-2017), Predicting Affordance Changes (Alternative ways to deal with uncertainty) School of Computer Science, University of Birmingham.
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html>

Sloman, A. (2013). Virtual machinery and evolution of mind (part 3) meta-morphogenesis: Evolution of information-processing machinery. In S. B. Cooper & J. van Leeuwen (Eds.), *Alan Turing - His Work and Impact* (p. 849-856) Amsterdam: Elsevier. (This was the first publication on the Meta-Morphogenesis project.)
<http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d>

Aaron Sloman (ongoing, under continual development since about 2012) *The Meta-Morphogenesis Project*
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

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

Sloman, A. (2013-2018). Jane Austen's concept of information (Not Claude Shannon's), Technical report. <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html>

Aaron Sloman, 2015, Some (Possibly) New Considerations Regarding Impossible Objects, (Their significance for mathematical cognition; current serious limitations of AI vision, learning and reasoning systems; and philosophy of mind--contents of consciousness). (Online research presentation,)
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

Aaron Sloman (2017). Construction Kits for Biological Evolution. In S. B. Cooper & M. I. Soskova (Eds.), *The Incomputable: Journeys Beyond the Turing Barrier* (pp. 237-292). Springer. (See next item, extending this.)
<http://www.springer.com/gb/book/9783319436678>

Aaron Sloman (kits). A revised extended version of [the previous item](#) is still under development here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>

Aaron Sloman, 2018, Alan Turing's 1938 thoughts on *intuition vs ingenuity* in mathematical reasoning
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html>
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.pdf>

[BACK TO CONTENTS](#)

Maintained by [Aaron Sloman](#)
[School of Computer Science](#)
[The University of Birmingham](#)