School of Computer Science University of Birmingham UK

Every intelligent ghost must contain a machine

an information-processing machine

Biologically Evolved Forms of Compositionality

Short version, for presentation at SYCO 1, 20-21 Sept 2018
Also available at https://goo.gl/DjDQNS

Aaron Sloman
School of Computer Science, University of Birmingham, UK
http://www.cs.bham.ac.uk/~axs
Last updated: 24 Sep 2018

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

---

Notes for a 25+5 minute talk on 20th Sept 2018, at
First Symposium on Compositional Structures (SYCO 1)
School of Computer Science, University of Birmingham, UK
http://events.cs.bham.ac.uk/syco/1/
(Includes schedule of talks)

These notes present a subset of the ideas in a longer (still growing) online paper:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.pdf
Also available as goo.gl/hQikAJ
Or: search for: Sloman and Compositionality

AIMS

This talk is mainly about compositionality in biology, with special emphasis on kinds of compositionality required for various types of spatial cognition.

I am also interested in how it came about that some products of evolution, namely, humans living several thousand years ago, included forms of spatial cognition that made possible the discoveries concerning geometry, made by ancient mathematicians, including those in Euclid's **Elements**.

Explaining those discovery processes and replicating them in AI systems is still beyond the state of the art in psychology, neuroscience, and AI.

In part that is because many researchers don't understand what needs to be replicated: e.g. it is not just the ability to get answers to mathematical questions right.

As Immanuel Kant pointed out in his *Critique of Pure Reason* (1781), mathematical discovery includes understanding why some things are impossible and some things are necessarily true, which cannot be done by collecting statistics and computing probabilities:
-- impossibility is not the same as 0% probability.
-- necessity is not the same as 100% probability.
How can a brain represent impossibility or necessity?

Ancient mathematicians were not using what are now called modal logics, and they were not reasoning about possible worlds:

their mathematical discoveries included facts about this world,
e.g. geometrical and topological facts.

A long term aim

Show how products of evolution that were selected because of their practical usefulness in many contexts, were eventually transformed into more general, less constrained, forms of information processing

including some forms of special-purpose spatial cognition shared with other species, which, in humans later provided foundations for deep mathematical discoveries.

Here’s a subset of key-ideas relevant to my topic:

Key Themes:
But not enough time for all:

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 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: a source of new kinds of entity -- including minds. Compare the relatively recent discoveries relating to types of virtual machine, including distributed virtual machines e.g. email systems, the World Wide Web.
Evolution made closely related discoveries long before we did.

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.

NOTE:
Frege implicitly introduced two parallel kinds of compositionality, insofar as he distinguished Sinn (sense, connotation) and Bedeutung (reference, denotation). This has important implications mentioned in the full paper.

A consequence that he played down is that in some contexts there can be a clear sense, yet a failure of denotation, e.g.

The father of the subject of this sentence is a mathematician

The subject of the sentence fails to refer, because of infinite recursion. So the sentence has no truth value: part of the information required to determine the truth value does not exist. There are many such cases which are often mis-diagnosed as generating contradictions instead of failure to have a truth value.
http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1971-03

Designers of programming languages normally try to make that impossible -- but in so doing they restrict the generative power of the languages.
END NOTE

METAPHYSICAL COMPOSITIONALITY

One of the features of biological evolution is metaphysical creativity: evolutionary processes can bring new kinds of entity into existence for the first time.

Examples include new kinds of information required to meet new control functions.

The forms of computation required for this included new kinds of virtual machinery, e.g. running on brains, or on insect social networks.

Compare the development by humans of new kinds of virtual machinery since the mid 20th century.

An example of this was the production of visual information processing.
At first computer scientists thought of visual mechanisms as pattern classification mechanisms: sensor signals were segmented into items to which labels were attached identifying recognized patterns.

Max Clowes and others replaced this with the idea of visual systems dealing "syntax" in the images received and "semantics" in the interpretations of those images.

Getting such systems to work included resolving ambiguities of grouping and using context to resolve ambiguities of reference (does this line feature represent a concave edge or a convex edge?)

This required perceptual systems to make use of different ontologies:
-- ontologies for sensory information
-- ontologies for depicted scene contents

In some cases that required multi-layered ontologies, e.g. when you see a face you may see not only the physical features but also an emotion or a mood, which involves using an ontology of mental states.

Unfortunately some of the advances made at that time have been undermined by recent work treating perception as if it amounted to multi-layered pattern recognition.

MODAL CONSEQUENCES

3D Visual spatial perception includes (at least) two levels of compositionality:
-- compositionality of 3D structure (scene syntax)
-- compositionality of 2D structure (projected image syntax)

The early AI researchers (e.g. Max Clowes, who introduced me to AI in 1969) pointed out that perfectly possible images/pictures could depict impossible scenes and therefore visual systems needed to be able to reason about impossibility. Many examples are here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html

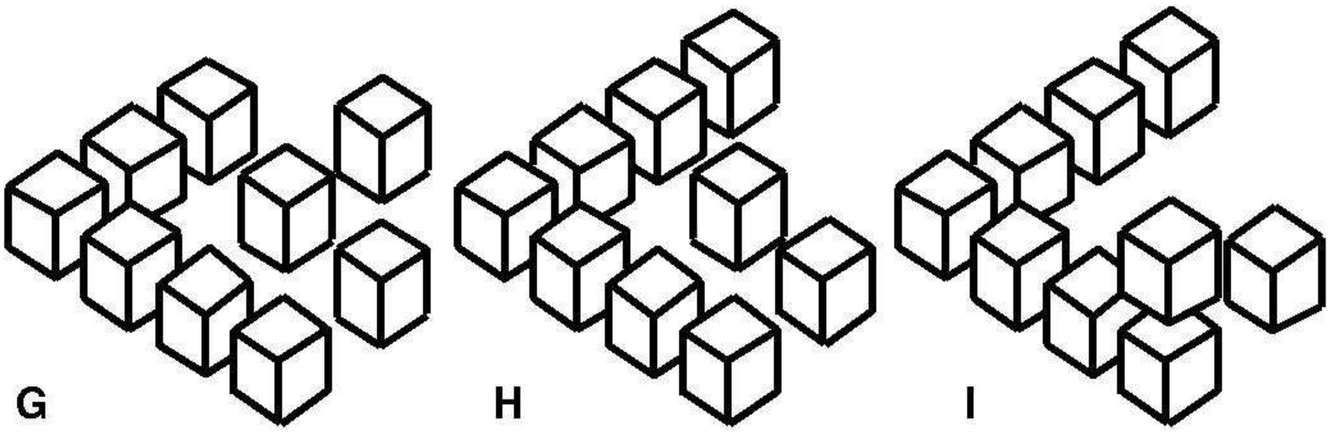

Here are nine blocks on a surface, shown in three possible configurations, including one in which one of the blocks is suspended above (or floats above) another block. How many other configurations are possible? How else could they be arranged? What sequences of block moves could produce the new arrangements?

Several more configurations of nine blocks are depicted below. You may or may not find one of them anomalous.

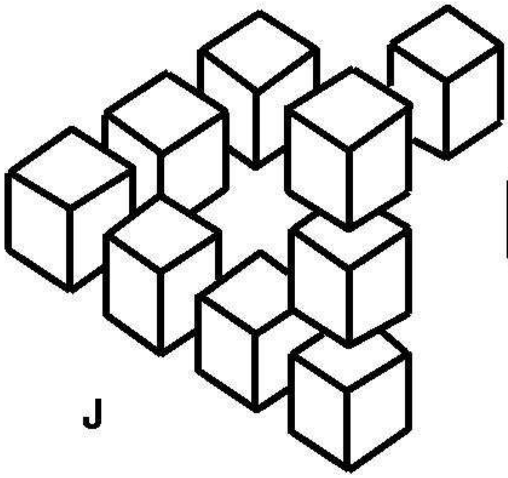
J

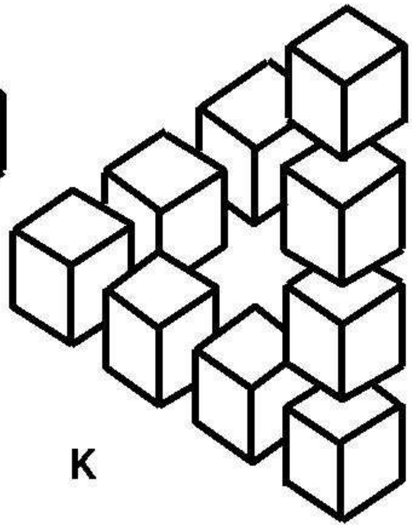
K

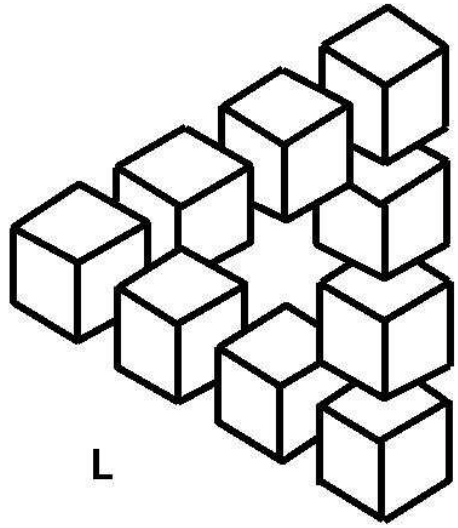
L

The above examples were all based on a picture by Oscar Reutersvard some time before the Penrose triangle.

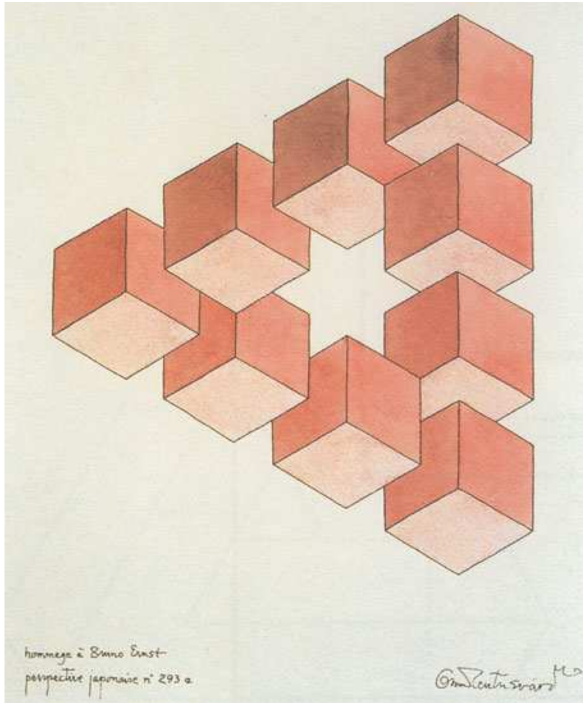

The ability to detect such impossibilities depends on abilities to reason about structural relationships, not the ability to generalise from masses of examples.

In this case, and many of Escher's examples inspired by the Penrose triangle, the impossibilities arise from composition of 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.

COMPOSITIONALITY IN EPIGENESIS

Waddington's "Epigenetic landscape"

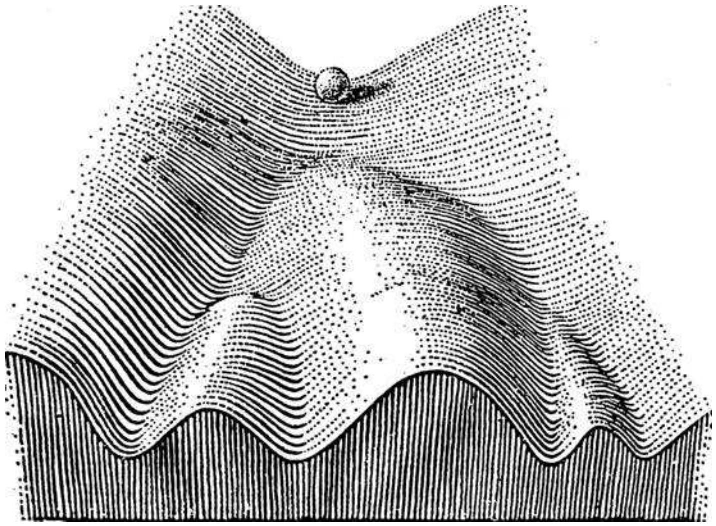

.

.

.

.

A META-CONFIGURED GENOME'S USE OF COMPOSITIONALITY

Upward and downward compositionality

Figure Meta

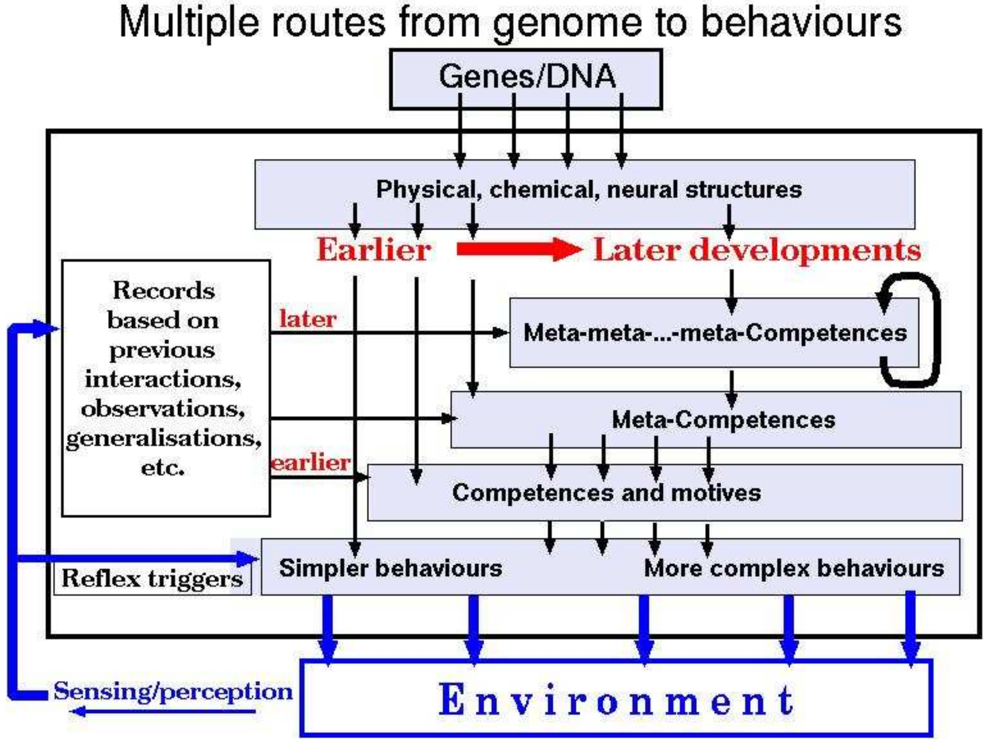

Figure Meta above, based on ideas in Chappell & Sloman(2007), shows (crudely) how staggered gene expression can allow relatively abstract/schematic results of later gene expression to be instantiated using information gained during earlier interactions with the environment -- not necessarily using any 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.

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.

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

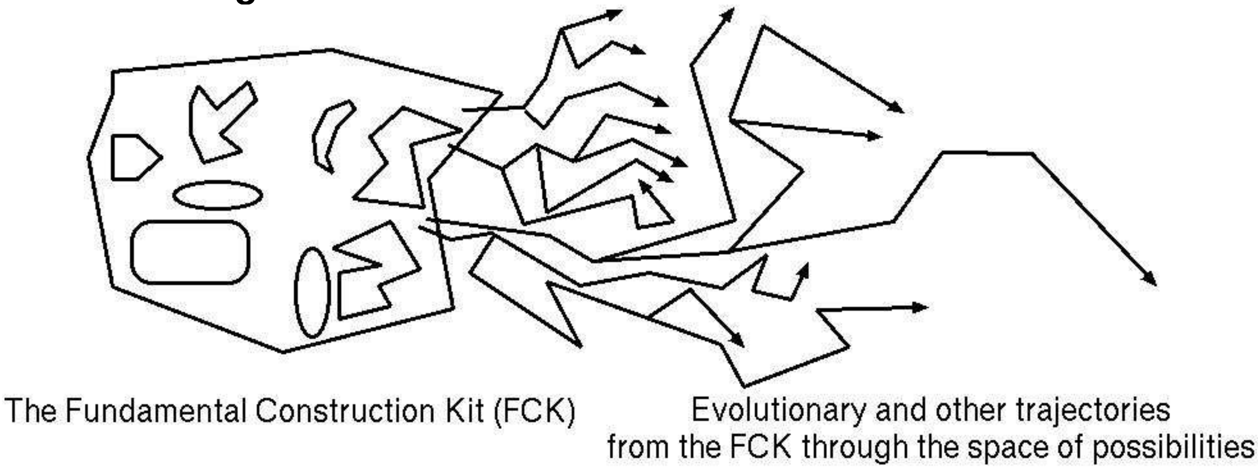

A crude representation of the Fundamental Construction Kit (FCK) (on left) and (on right) a collection of trajectories from the FCK through the space of possible trajectories to increasingly complex mechanisms.

Figure 2 DCK: Derived Construction Kits

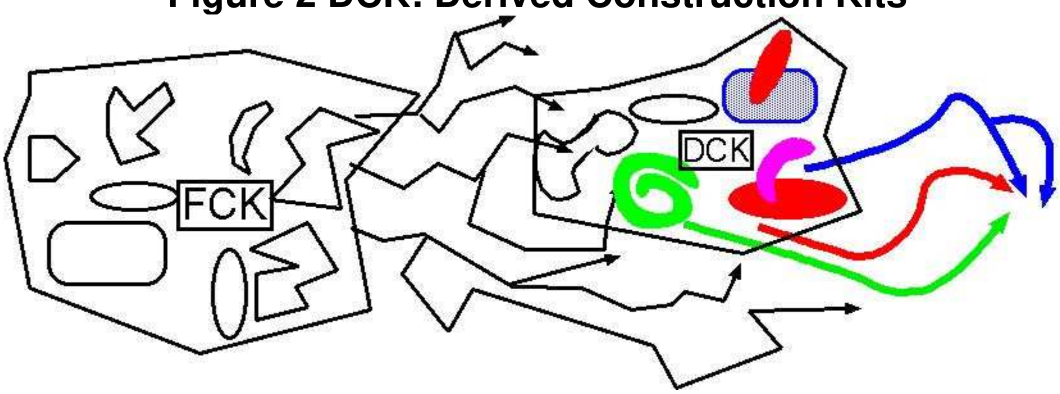

Further transitions: a fundamental construction kit (FCK) on left gives rise to new evolved "derived" construction kits, such as the DCK on the right, from which new trajectories can begin, rapidly producing new more complex designs, e.g. organisms with new morphologies and new information processing 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 *Fundamental* Construction Kit (FCK) provided by the physical universe made possible all the forms of life that have so far evolved on earth, and also possible but still unrealised forms of life, in possible types of physical environment. Fig. 1 shows (crudely) how a common initial construction kit can generate many possible trajectories, in which components of the kit are assembled to produce new instances (living or non-living).

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

ADDITIONAL TOPICS IN THE PAPER

Evolution of ancient mathematical abilities

Compositionality in natural language

Intensional and extensional compositionality

Reasoning about appearance changes caused by motion

Spider creativity

More to be added!

Huge gaps between current AI and the forms of intelligence produced by biological evolution.

Evolution could not produce its intelligent animals by training them on data: the data did not exist!

I suspect our current models of computation have some deep unnoticed limitations, though it looks as if Alan Turing began to explore ways of filling the gap shortly before he died, in his 1952 paper:
The chemical basis of morphogenesis.

What would he have done if he had lived several decades longer?

My guess is: The Meta-Morphogenesis project
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html

I've provided some snapshots of its contents. Much help needed.

Can category theory -- a major theme at the workshop -- help?

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