



The Creative Universe Original title: Creativity and AI

(INCOMPLETE DRAFT: STILL UNDER CONSTRUCTION)

Aaron Sloman
School of Computer Science, University of Birmingham

Installed: 4 Mar 2016

Last updated: 16 Mar 2016; 2 Dec 2017; 29 May 2018; 22 May 2019

This paper is available as html and pdf

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

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

Some of the key ideas were in Chapter 2 of The Computer Revolution in Philosophy (1978), presenting the aims of science as discovery and explanation of **possibilities**, including discovery of generative explanatory mechanisms, rather than discovery of **laws** and law-based explanation.

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

So key questions include: How is this possible? What makes this possible? Discovery and explanation of laws is secondary to discovery and explanation of possibilities, insofar as laws are merely limitations of possibilities.

This is now part of the Turing-inspired Meta-Morphogenesis (M-M) project, begun late 2011:

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

A partial index of discussion notes is in

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

CONTENTS (Likely to change)

- [Varieties of, and contexts for, creativity](#)
- [Architecture-based vs Reward-based motivation](#)
- [Not only intention-driven creativity](#)
- [Perceptual creativity](#)
 - [Note on mathematical creativity](#)
- [Creativity depends on uncreative mechanisms](#)
- [Deep structural creativity](#)
- [Mind-less \(not mindless\) creativity](#)
- [Back to what evolution can do.](#)
- [Unclear beginnings](#)
- [Evolution's creativity](#)
- [Stuart Wray's sketch of meta-morphogenesis](#)
- [From planet-wide creativity to individual creativity](#)
- [Creativity is polymorphous](#)

- [How did evolution produce mathematicians like Euclid and his precursors?](#)
- [Requirements for new biological states and processes](#)
- [Biological creativity increases spontaneously](#)
- [Creativity in Artificial Intelligence](#)
- [Varieties of computational creativity research](#)
- [Is it possible for computers to be creative?](#)
- [Types of post-boot behaviour](#)

[More links to contents to be added]

- [REFERENCES AND LINKS](#)
-

Varieties of, and contexts for, creativity

It is normally assumed that only an individual acting intentionally can be creative. I'll try to show that that is a special case. There is a deeper notion of creativity that is applicable to any kind of causal agent or mechanism -- whether intentional or not, alive or not -- that has available a class of possibilities for action or change. Anything done by that agent must already have been included in that class of possibilities. In that sense, products of creativity are not entirely new: what's new is realisation of a pre-existing possibility.

However, as we'll see, in some cases the pre-existing possibility is implemented in a distributed form: various things that support that possibility may never previously have come together in this universe (or merely on this planet?). In that sense some possibilities can't be realised without first realising some intermediate possibilities -- e.g. creating instances of some subset of possible tools or construction kits, or conditions for realisation.

In that sense the universe supported the possibility of carbon based life and many more specific possibilities for forms of life and their products long before any parts of the universe had the right combination of physical mechanisms to underpin the [realisation](#) of that possibility.

A brief summary of the main ideas I'll present goes as follows: Creative processes or mechanisms and their products can be compared in various ways, e.g. on the basis of the complexity of the [process](#) of realisation or complexity of the [result](#), the [differences](#) between the new product and previously realised products, the relative [likelihoods](#) of realisation of various alternative possibilities or sets of possibilities, changes in "[step size](#)" of actions or opportunities, e.g. whether features of the new range of possibilities are made [directly](#) accessible, the size and variety of the new range of possibilities, and, in some cases, the [possibilities made inaccessible](#) after realisation of new certain possibilities (e.g. because of resources consumed, opportunities rejected, things destroyed, or other incompatibilities), and other criteria that can vary according to the type of product, and also the interests or needs of whoever or whatever is doing the comparison.

Realising a new possibility may involve acting intentionally or unintentionally. Natural selection is a process that makes changes unintentionally -- at first, until one species starts breeding another, or other more subtle processes evolve, e.g. uses of new cognitive powers in mate selection. Some of the greatest creative steps taken by evolution occurred before any of its products were capable of acting intentionally, for example the evolution of mechanisms of sexual reproduction.

Much of evolution involves **blind** creativity: nothing is aware of options available, options chosen or consequences of the choices. This is unlike a "blind watchmaker" who has goals, preferences, and abilities to recognize improvements, but cannot perceive opportunities, materials, or design possibilities, in advance of their use.

Another example of blind creativity is the collection of physical mechanisms that determine the path of an electronic discharge during a thunderstorm. Every path may be unique, but the uniqueness is neither recognized nor of any value or benefit to the physical system producing the discharge.

In some cases realisation of a new possibility, such as construction of a dam by beavers, depends on a mixture of pre-existing physical resources and constraints and the cognitive processes controlling the physical behaviours of the individual beavers.

In general, it is not possible to tell simply by observing behaviours whether there is a prior intention to produce the result that appears to show intentional creativity. For example, beavers seem to have the intention to create dams that raise water level, but an experiment has been reported that suggests that the intention of each individual beaver is only to achieve reduction of noise from fast running water or (whatever noise source triggers the dam-building behaviour), at least insofar as noise reduction is what leads to termination of the behaviour:

http://www.naturealmanac.com/archive/beaver_dams/beaver_dams.html

Architecture-based vs Reward-based motivation

Similarly, behaviours in young children that appear to be based on intentions to produce some future benefit that the behaviours actually do produce may not be based on any such intention, but on a causal mechanism previously selected by evolution because the individuals that had that mechanism survived long enough to have offspring and actually had offspring with the same genetic predisposition. The motivations are **architecture-based** rather than **reward-based**.

Examples are discussed in [Chappell and Sloman \(2007\)](#), and in [Sloman 2009](#)

For example a great deal of activity goes on in a normal pre-verbal human infant or toddler that contributes to development of linguistic competences later on. A linguist watching may know that, but not the child. Neither the future benefit nor expectation of future benefit plays a causal role in the operation of the mechanisms. Some of the ideas (and the main diagram) in the 2007 paper have since been extended, as briefly summarised in

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

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html#tth_sEc2.3

In some cases the mechanism is a reflex action that is triggered by the environment, e.g. a blinking reflex when something moves rapidly towards an eye. In other cases the reflex doesn't produce a bodily action, but creates a goal which may or may not be selected for action, depending on what other goals compete with it. If the goal is selected it may provide information about the environment that is stored, and may or may not contribute to some important later development (e.g. formation of a theory about the world, or formation of a grammar for the local language).

I call that "Architecture-based Motivation" (ABM) in contrast with Reward-based Motivation (RBM), discussed in [Sloman \(2009\)](#). The mechanism seems to be very clearly at work in spontaneous playful activities of the young of humans and other species, which are often highly creative insofar as individuals do things they have never seen done by others and have not found previously to be

useful. From the standpoint of this paper these are merely particularly sophisticated recent manifestations of types of creativity ultimately produced by mechanisms that originally existed in a lifeless universe. But less obvious forms of playful creativity seem to abound in the proliferation of forms of life.

Often scientists studying such actions assume that if the child is unaware of the main long term benefit then it must perform the actions for the sake of some immediate reward (some sort of positive reinforcement provided by brain mechanisms). But that is simply an unjustified assumption because they have not had the experience of designing mechanisms that are capable of inserting goals into a decision-making or planning mechanism without also inserting some promise of a reward, or an actual reward if the goal is achieved. So those theorists are seriously lacking in scientific creativity.

Although we normally think of creativity as a feature of intelligent agents, acting intelligently, with an intention to achieve a goal, solve a practical problem, answer a question, find out what will happen if...., or to create something intrinsically interesting or worthwhile, such as a work of art, an artistic performance, or a form of play, I have shown that there can be creative mechanisms at work without any of those characteristics.

Not only intention-driven creativity

Intention-driven creativity can occur in a child, in a squirrel, in a community solving a problem, in an engineer, in an engineering company, in a research community, and many more. But that that's not all there is to creativity.

The most familiar varieties (and most written about varieties?) of creativity that I know of involve intentional, intelligent action. There are many examples of animal minds or groups of minds finding a way to solve a problem that it or they previously could not solve. The problem may or may not have been identified as a problem earlier. Sometimes the new method of solution is stumbled across, by accident, but recognised and used, perhaps because all other alternatives considered have failed to produce a desired effect. It is also possible for a creative action to be *unintentional*: e.g. something done by accident, can turn out to be a solution to a previously unsolved problem.

This paper is about what makes something creative, whether it is done intentionally, or consciously or not. It need not even be produced by intelligent agents. Many of the changes produced by biological evolution are highly creative in that sense, including large scale changes produced by many relatively uncreative small changes.

Intentional creative actions are often the result of an individual studying a problem situation or a mechanism of some kind and noticing or working out that a novel strategy will solve a new hard problem involving that situation or mechanism. For example, a squirrel that is at first defeated by a rotating bird feeder designed to hurl visitors heavier than birds away, may notice that nuts are thrown out during the rotation. If it then returns to the feeder and hangs on for several rotations, while nuts are thrown to the ground, and then drops to the ground to eat the nuts, it has creatively defeated the anti-squirrel mechanism. (Diligent searching will reveal online videos illustrating such behaviours.) Creative designers of anti-squirrel mechanisms may respond by modifying the mechanism to prevent nuts being ejected by centrifugal forces.

Creative processes can occur in a game-playing machine, if it is designed not merely to play the game, but also to monitor and record features of games it plays, looking for clues regarding good and bad moves. As early as 1959, Art Samuel had designed a self-teaching checkers program that was able eventually to defeat its author. Details are available in [Samuel\(1959\)](#). Although a human can intervene if the program does not seem to be learning as expected, Samuel states: *In general, however, the program seeks to extricate itself from traps and to improve more or less continuously.*

Perceptual creativity

Perceptual processes can also be creative, as I learnt from Max Clowes around 1969, when he arrived at Sussex university, to continue his AI research in vision, while I was a young lecturer in philosophy. I learnt from him that the task of a vision system making sense of a picture, even a photograph of a simple and familiar type of scene, could involve forms of creativity that were not present in the most common kinds of "vision" programs (at the time), which generally used pattern recognition techniques to attach previously learnt labels to an image, or parts of an image.

His examples involved not just recognition or labelling of images or image parts, but something more complex, involving perception of parts, relationships between parts, ways of grouping parts, ways of "making sense" of some grouped parts (e.g. seeing them as parts of a partly obscured object, represented by disconnected parts of the image, even though the object is fully connected).

Many cases require visual mechanisms to infer (creatively) that an image structure not seen previously is the projection of a familiar object from a new viewpoint. In some cases the results include noticing possibilities for changes that will produce new situations, with new properties. E.g. cutting a small piece off a large lump of food can produce something small enough to fit into one's mouth: a new possibility.

In all these cases the processes of vision produce something new that is structurally related to the visual input, but the relationship need not involve simple mappings between image structures and perceived objects or scenes.

Creative discovery in mathematics may involve perceiving new possibilities for alteration of a previously known type of structure, or consequences of such possibilities, or constraints on such possibilities.

Note on mathematical creativity (21 Feb 2019)

For a more detailed discussion of evolution's mathematical creativity (discovering, combining and using mathematical structures) see

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-framephys.html> (or [pdf](#))

This is closely related to the roles of [compositionality](#) in biological evolution often wrongly thought to be merely a feature of language use:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html> (or [pdf](#))

Even when something is labelled using a previously known category, the perception process will produce a lot more than the label: namely a collection of parts and relationships of things seen and also relationships between parts, e.g. X's left arm is touching a leg of the table. Often what is seen *explains* what is experienced in the details of the image.

NOTE

This is part of an argument I first used in 1978 supporting the claim that evolution must have produced languages for *internal* use in organisms long before languages were used for *communication*, as explained in

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

I have noticed that many researchers have a conception of "language" that is so rigidly (uncreatively) linked to communication between individuals that they are incapable of thinking about any alternative function for human language. That applies even to people who are familiar with types of programming language, or data-structure specification languages, for example, used internally in running computing systems.

An exception is the mathematician, [David Mumford](#) who reached similar conclusions independently.

Before coming to Sussex Max Clowes had written about the creativity required for visual mechanisms interpreting images or diagrams. Examples are provided in the obituary notice extended with bibliography and comments [Sloman \(2014\)](#). His approach contrasted with work on statistical mechanisms that could train a computer to recognize patterns in images: he stressed the importance of seeing not the *category or contents of the image*, e.g. a 2-D structure on a page, but *other structures* (e.g. a 3-D structure, or an electronic circuit) *depicted in the image*. Later he wrote about "Man the creative machine", partly as a result of working on requirements for human-like visual capabilities in machines [Clowes \(1973\)](#).

Of course, the operation of the learning mechanisms in Samuel's program (and other learning programs) resulted from earlier decisions by a human programmer. In the following decades, with far more computer power than was available in 1959, computer programs could do far more sophisticated exploration and discovery at different levels of abstraction, including machines playing far more complex games, such as Chess and Go. Such systems creatively extend their competences in part by using biologically inspired forms of computation, including "evolutionary" mechanisms and multi-layer "neural net" mechanisms inspired by (possibly incorrect) theories of brain function, as well as other techniques developed in AI and software engineering (for instance engineering techniques required to make hundreds of computers cooperate in playing against a human, or managing a flight control service).

NOTE

I think all such programs still lack (in 2018) biologically important visual and spatial reasoning capabilities of kinds that led to human discoveries in geometry and topology eventually brought together in Euclid's [Elements](#). A subset of the abilities required seem to be shared with other intelligent animals, and pre-verbal children. Further discussion of the required mechanisms can be found in

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

Creativity depends on uncreative mechanisms

Creativity always includes unintentional components, such as inherited information processing mechanisms. Even a young human poetic, musical, mathematical, engineering or sporting prodigy is not completely self-created: the language used and many intellectual and physical competences will partly have been absorbed from the culture, and some of the computational mechanisms and learning strategies used must have been discovered earlier by natural selection in the individual's

past. Originally no organism had them.

Biological evolution and patterns of development, learning, and creative activity in individuals, show that the physical universe can support processes of evolution, and processes of learning, that over time produce more and more types and levels of novelty, without any external intervention, and without any need for all the changes to be results of previously formulated intentions, whether human or non-human.

So a general theory of creativity must not assume that creativity necessarily involves solving a previously recognised problem, meeting a previously recognised need, or fulfilling a pre-existing intention.

It doesn't even require the existence of a mind that is creative: creativity of the type under consideration here may simply be an inherent feature of a process or mechanism, that creates novel instances in a space of possibilities, or creates modifications to one or more spaces of possibilities that generate a new space of possibilities, or combines two or more possibility spaces to produce previously unreachable types of complexity, or discovers re-usable features of a space of possibilities, e.g. by creating detectors for members of a certain subset of the space, and making use of the detection either by avoiding, or by seeking and deploying those newly characterised possibilities (e.g. avoiding detected weaknesses in a construction process, or deploying specific constructions that yield new desirable features).

Biological evolution uses such mechanisms. (An exercise for the reader would be to identify and analyse examples of all these types on different scales in biology, e.g. molecular scales, individual organism scales, community scales, ecosystem scales, etc.)

In particular, evolution creates new construction kits that change what evolution and its products can bring about and can prevent, and until now, in the variety, range of physical scales (space and time), types of material, types of function, types of automaticity, and even types of information processing it has far surpassed the creativity of human designers, though humans may produce some new designs much more quickly: using resources indirectly provided by biological evolution (working through brains of human scientists and engineers).

In other words, the physical universe supports a type of biological evolution that exhibits deep forms of compositionality, as discussed in

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

Deep structural creativity

If we want to understand the varieties of creativity in this universe we can focus on production of novelty, whether it happens quickly or slowly, and whether it involves intentions of agents or not. We'll find that great creativity does not always imply the pre-existence of a mind with intentions, goals, beliefs, strategies, etc., directing or controlling the processes.

Evolutionary creativity is a feature (I would call it a *structural* or *mathematical* feature) of the mechanisms and processes of evolution by natural selection, supported by features of the physical universe, some of which were described by Schrödinger, in his [1944 book](#), in which he explained how quantum mechanisms make it possible for complex molecules composed of lengthy aperiodic chains of atoms to form a richly expressive and highly stable medium for preserving genetic

information across many generations. (This was published before Shannon's ground-breaking paper, anticipating some of Shannon's ideas -- but see note on Shannon's misuse of the word "information" [below](#).)

Schrödinger's ideas in this little book also help to explain the notion of a Fundamental Construction Kit that supports all the later construction kits produced by natural selection. The combination of structural stability and controlled switchability of molecular structures that he identified are a deep part of the explanation of biological creativity.

Deep creativity existed in the mechanisms underlying biological evolution (the biological "construction kits") long before there were any human or animal minds. The products of that creativity included new, even more sophisticated, forms of creativity in later construction kits and, still later, construction-kit users, with increasingly complex minds.

Those varieties of creativity, are inherent in physical and biological processes that make a difference between what exists at one time and what exists later on.

For example, if our planet was once a collection of chemical substances held together by gravitational attraction supplemented by chemical bonds, orbiting around a sun, without any forms of life, and was able to transform itself (albeit using external sources of energy), to what we now have, with millions of enormously diverse life forms, a subset of which can discuss evolution and creativity, then the universe that produced and supported that planet without the intervention of any pre-existing intelligence, is clearly deeply creative.

Mind-less (not mindless) creativity

In that sense, creativity, normally considered to be a feature of [minds](#), actually has nothing specific to do with minds, but is a general feature of some combinations of processes and of mechanisms producing further creative processes.

Production of minds (and their products) is a subset of a more general class of processes that I am trying to characterise, for which certain deep features of the physical universe seem to suffice.

(I believe a partly similar thesis is presented, with detailed supporting evidence in Stuart Kauffman's [\(1993\)](#) book which I have only looked at (briefly) for the first time a few days ago, though I had previously read his [\(1995\)](#) book, whose influence on my thinking may be deeper than I know. More recently, related ideas have been expressed by [David Deutsch \(2011\)](#), though without explicit discussion of layered construction kits.)

A separate paper on the roles of construction kits in biological evolution

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

distinguishes varieties of construction kit, starting with the Fundamental Construction Kit (FCK) provided by the physical universe, using features noted in [Schrödinger \(1944\)](#), and others.

Various physical, chemical and biological processes produced layer upon layer of new Derived Construction Kits (DCKs) of various sorts, including concrete, abstract and hybrid construction kits [Sloman \(2017--\)](#). All of the construction kits extended the creativity of (some part of) the universe, by shortening the times required for reaching various new types of life, including physical forms, new physical behaviours and new forms of information processing, and increasing the diversity of possible products.

In some cases, the new mechanisms, or their products, removed some of the previous generative potential in neighbouring parts of the universe, by changing the environment in ways that prevented continuation of some previously enabled processes. E.g. production of an oxygen-rich atmosphere helped some organisms and impeded others.

I believe this is related to but richer than the notion of an extremely "rugged" fitness landscape discussed by Kauffman and others, since we are talking about "interacting fitness landscapes" for different types of organism. E.g. it is possible that conditions that allow some structures, e.g. porous rocks, to grow without developing a life-form might provide "scaffolding" for a new life form that makes use of those structures, for food, for shelter, or temporary support during development of a self-supporting body.

Our ability to make sense of this was enormously expanded during the last five decades of the last century during which we discovered for the first time how to make physical machines that could support many different sorts of virtual machinery, interacting with one another and with physical parts of their environment. This had been done much earlier by biological evolution in the process of producing many varieties of information processing systems including animal minds. (I tried to make this point in [Sloman \(1978\)](#), but failed to communicate with most readers, especially in [Chapter 2](#).)

Back to what evolution can do.

Long before there were organisms that could be described in terms of perception, learning, knowledge, intention, decision-taking, or goal-seeking, evolution was producing spectacular, changes in designs and functionality of organisms of many types and sizes, from molecular scales to giant trees and huge sub-soil fungi^[*], along with even larger intricately functioning ecosystems. Those are examples of evolution's creativity without intention, and without awareness of products or their effects.

[*] <http://www.scientificamerican.com/article/strange-but-true-largest-organism-is-fungus/>

This is different from a mindless rearrangement of matter that has many consequences, such as a volcanic eruption, a major earthquake, an asteroid impact, or the initial formation of this galaxy, the solar system, or the earth. One difference is that the products of evolutionary creativity are, or are functioning parts of, complex, organised, self-maintaining reproducible mechanisms.

Some features of such self-organising feeding, growing, structures, were assembled as "toy" creatures named "Droguli" (singular "Drogulus") constructed from wood, screws, springs, hinges, etc. and demonstrated "feeding" and reproducing themselves in a lecture I attended in 1959 by Lionel Penrose (father of Roger):

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

At a still unknown stage, the physical universe somehow produced the very simplest self-maintaining, reproducing organisms, perhaps with the Chemotron mechanisms proposed by [Ganti \(1971/2003\)](#) as meeting minimal requirements for life). Even producing such a "minimal" life form required a deep kind of creativity in the physical universe, since the physical and chemical structures and mechanisms available initially on the planet were far simpler.

It is likely that non-biological, non-evolving, physical and chemical processes produced significant substructures that could function as self-maintaining proto-life forms before the full functional complexity of a Chemoton existed.

Unclear beginnings

Before the existence of the most primitive life forms there were physical and chemical structures of varying complexity (from sub-atomic particles to galaxies, etc.). These could have included spontaneously produced increasingly complex molecules capable of playing roles in living organisms or providing environments in which life forms could be assembled and supported. [Schrödinger \(1944\)](#) showed how quantum physics made possible the formation of mechanisms required for (long term) storage and use of genetic information needed for reliable biological reproduction.

The same features of the physical universe also allowed the formation of increasingly complex stable structures, resistant to thermal buffeting and other disturbances, which, in turn, made possible yet more structures formed either by combining them, or by using them as temporary scaffolds during assembly of even more complex structures, including continents.

Not all stable, self-maintaining structures are manifestations of the creativity required for life. For example the earth plus moon system exhibits dynamic stability and has regular, predictable effects, such as tidal flows. But such phenomena are fully explained by theories in physics, and do not exhibit any ability to acquire information about their current needs and available opportunities and redirect forces to make use of those opportunities -- unlike even the simplest organisms selecting molecules at their surface, excreting waste products, and using some absorbed chemicals as food.

Such an organism, however simple, uses some of its own energy to prevent, reduce the impact of, or terminate harmful processes; or to initiate useful processes e.g. finding food, ingesting food, resting, acquiring information, etc. While growing itself, maintaining itself, repairing damaged parts of itself an organism uses available or previously stored information to generate alternative solutions and to choose one, unlike a storm cloud, an avalanche, a drifting continent or a planet.

Not only encoded information, but also robustly functioning systems require mechanisms that allow both a great deal of variability and a great deal of stability. (I think Schrödinger was saying something like this in the final chapter of his little book, but I am not sure.)

Evolution's creativity

After the simplest life forms, there were spectacular changes in the types of self-organising, self maintaining, self improving processes, supported by changes in mechanisms producing those processes. Those processes include the operations of increasingly sophisticated physical/chemical construction kits, and increasingly sophisticated information-processing mechanisms produced by evolution: the main topic of the Meta-Morphogenesis project.

At first the mechanisms changed only as a result of blind (random) variation and natural selection, supported only by the sorts of potential in physical/chemical matter identified by Schrödinger. But later, increasingly complex goal directed learning mechanisms helped to speed up change, and steadily produced more complex results. Richer forms of information processing could support the preservation and proliferation of organisms with more complex and varied nutritional needs.

Evolution's use of increasingly complex forms of compositionality, including use of "meta-configured" genomes, helped to accelerate the rate of change, or at least made bigger "steps" possible, as explained in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/compositionality.html> (also [PDF](#)) and <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html> (also [PDF](#)).

Spectacular examples of the increasing rate of change of products of evolution are included in the recent changes in human culture resulting from a constantly accelerating collection of interacting engineering design developments since World War II -- one side effect of which was the production of the internet, followed later by the World Wide Web, and a rapidly increasing variety of new mechanisms and applications based on that. These are ALL products of biological evolution, insofar as they are products of products of biological evolution. ("Product of" is a transitive relation, unlike "Direct product of".)

Summary so far:

Deep features of the physical universe made it possible for naturally occurring processes to assemble not only increasingly complex structures and processes, but increasingly complex novel "platforms" for producing things that could not previously have been created. So things that were impossible when our planet first formed became possible much later. In some sense, of course, they were always possible, but possibilities that were initially unreachable became reachable. Or, in other words, possibilities that required enormous numbers of steps to realise them, later could be realised in far fewer steps. Some of the transitions making certain life forms possible may also have made others impossible: for example life forms that depend on a very different atmospheric composition are impossible in our atmosphere, in which the amount of oxygen is one of the many products of biological evolution, including evolution of plants capable of photosynthesis.

In that sense the price of use of creativity is sometimes blocking of creativity. Other examples are likely to exist in the development of brains initially capable of realising far more possibilities than any one brain ever realises. In the process of developing certain linguistic and cultural competences human individuals may, like evolutionary processes on our planet, exclude other possibilities that were initially open to them.


Stuart Wray's sketch of meta-morphogenesis (indicating some of the creativity of evolution) (copied from the M-M overview document).

On 5th Jun 2012, [Stuart Wray](#), after reading a draft conference paper on Meta-morphogenesis and the Creativity of Evolution (before the ideas about construction kits had been added):

<http://www.cs.bham.ac.uk/research/projects/cogaff//12.html#1203>

produced a sketch of the ideas in the project, reproduced here, with his permission:

(CLICK THE PICTURE TO SEE A LARGER VERSION)



[CLICK HERE FOR PDF VERSION.](#)

From planet-wide creativity to individual creativity

The importance of creativity in AI was recognized by some of the earliest AI researchers, including Marvin Minsky summarising progress and possible future steps in AI in a survey paper as long ago as 1960 [Minsky \(1960\)](#).

Often evolution is contrasted with creation, where "creation" refers to some sort of god. I have argued [elsewhere](#) that claims about the existence of a god in the most influential religions are neither true nor false, but incoherent. I labelled that thesis "Analytical Atheism" (a label also apparently favoured by a Facebook group <https://www.facebook.com/anatheism/>). I shall therefore not waste time discussing anti-evolutionary creationism, apart from noting the incoherence of theistic versions.

However, it is not incoherent to claim that evolution involves creativity since the processes of evolution by natural selection have clearly demonstrated enormous creativity, including all forms of human creativity since they depend on the operation of products of evolution.

The creativity of evolution uses quite different time-scales from human creativity, but that does not rule out deep structural commonalities.

The claims I am making contrast with the assumption that creativity requires something like consciousness and intentionality, or at least the use of a mind. That seems to be assumed by most researchers on human creativity, including one with whom I agree on many other topics. Boden [\(2004\)](#), states:

"But only the mind can change itself. And only the mind changes itself in selective, intelligible, ways. The 'journey through musical space' whose travellers included Bach, Brahms, Debussy, and Schoenberg was a journey which not only explored the relevant space but created it, too. In short, only the mind can change the impossible into the possible, transforming computational 'cannots' into computational 'cans'." (p61)

However she comes very close to the more general claim that I am making, when she writes:

"Biological evolution has had many millions of years in which to generate a wide variety of novelties, and in which to weed out the useless ones. But we must improve our thinking within a single lifetime (or, collectively, within the history of a certain culture or the cumulative experience of the human race)." (p226)

This implies that biological evolution is capable of great creativity, including the ability to judge or compare its products, and select a subset.

This illustrates the point that creativity does not require intentionality, since the physical universe, which includes all the mechanisms and products of biological evolution, is clearly the most creative thing that exists, being the **unaided** source of all forms of life, mind, evolutionary creativity, animal creativity, cultural creativity, ecosystem creativity, etc.

How it manages to be so creative needs a deep explanatory theory, ideally an expanded much better developed theory of construction-kits than the one being developed here. What is obvious is that the physical/chemical universe provides a "giant" construction kit, that includes, among other things, the ability to bootstrap new construction kits (new concrete, abstract and hybrid kits, including "temporary scaffolding" of various sorts, etc.), without which no form of life, on this or any other planet. would exist.

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

Boden's ideas of P-creativity and H-creativity (personal and historical creativity) [REF] transfer without much modification to various products of evolution, at various levels of complexity. Obviously there is also what she calls "transformational creativity" throughout evolution, because of the number and variety of evolutionary developments that change what becomes possible among living individuals and ecosystems on the planet.

One of its by-products on this planet is cultural evolution. A spectacular example, that stands out among the products of human engineering, occurred only since the 1950s, with steadily increasing complexity, sophistication of powers, and variety of applications, namely the development of increasingly complex human-designed **virtual** machines, with increasingly complex and indirect relationships to the underlying -- and surrounding -- physical machinery.

What is not widely recognised is that these human achievements depend on mechanisms that create and use virtual machinery produced by evolution and used in the minds of intelligent organisms long before the development of (much simpler) virtual machines by human engineers.

These were produced by some of the most recent products of evolution, including human computer scientists and engineers! A partial survey of the variety of types of information-processing virtual machines designed by humans is presented here:

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

Most philosophical theories of emergence, of supervenience, of ontological variety, fail to do justice to all that complexity -- mainly because learning how to think about such things (partly by designing, building, testing, debugging and extending working examples) is not yet part of a standard philosophy syllabus, as it should be, including detailed study of AI mechanisms and how to build, debug, modify and extend them. (My 1978 prediction that this would become a standard part of philosophical education within a decade or two proved disastrously wrong.)

But long before these amazing and powerful new products of human creativity, biological evolution had exhibited even greater creativity, e.g. in producing so many minds in so many species of intelligent animal, with far more varieties of virtual machinery than anything so far produced by human engineers.

Some of the most complex biological virtual machines that are still not understood are involved in visual perception. Many years ago Max Clowes, who introduced me to programming and Artificial Intelligence when he came to Sussex University taught me (and others) that many familiar processes of visual perception in humans require creative production of interpretations (e.g. 3-D structures and processes) of features of images using processes that can be based on simple generic mechanisms, that work for all images and scenes. See his 1969 paper briefly summarised here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html#mbc-1969>

Much of my own work on vision in AI is built on, and extends those ideas. But I have mainly been assembling requirements rather than designing and building solutions, e.g.:

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

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

Although Clowes did not point this out, the requirements for the kinds of creativity that he discussed in human visual perception must, to a considerable extent, apply also to visual processes in other animals, e.g. a crow looking for a suitable place to insert a new twig in a partly built nest. (There has been a lot of work on creative problem solving in crows and other corvids. [REFS])

Creativity is polymorphous

The claim that creativity (e.g. in biological evolution) does not require intentionality whereas other forms of creativity do, is related to the fact that "creativity" or "being creative" is a **polymorphous** concept. That is creativity is a polymorphous mathematical property of combinations of structures and processes and their causal powers. I shall later try explain what that means, with examples, but the key ideas are in this paper, which distinguishes various kinds of conceptual complexity that are often mistaken for vagueness, or "family resemblance", or statistical concepts:

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

Many concepts that I used to think are "cluster concepts" (i.e. loosely constrained disjunctions of more specific concepts, something like Wittgenstein's notion of a "family resemblance" concept) are polymorphic in that sense, including "conscious of/that", "efficient", "explains", ...

The first example of this kind that I noticed was "better", analysed in a paper published in 1969 (after a study of **Which?** consumer reports) but I did not use the label "polymorphous" until much later, after encountering precise and varied but very powerful examples in the design of advanced programming languages.

My mathematical specification in 1969 was very clumsy (despite help from a distinguished mathematician!).

In retrospect I think that [Chapter 2](#) of [Sloman \(1978\)](#), was claiming that 'science' is also a polymorphous concept concerned primarily with creative processes that generate theories about what's **possible**, and then intersperse them with 'laws' restricting possibilities, i.e. causally necessary connections. (The formulation of those ideas still needs a lot of improvement.)

Gilbert Ryle seems to have understood the importance of polymorphism in 1949, in **The Concept of Mind** but at that time he could not have grasped the full richness/precision of the concept that later emerged in mathematics and computer science.

How did evolution produce mathematicians like Euclid and his precursors?

The main kind of creativity that I am struggling with is what explains the evolution of mathematicians like Euclid, and the many unsung predecessors that must have contributed along the way. It requires invention of at least two, and probably three or more layers of meta-cognition, one to represent possible, but not actual variations of actual configurations, and one to detect (and use) constraints on those possibilities (which I think intelligent species do) and another layer to notice that all that is going on and think about particular cases, teach them to others, argue about them, and then another **systematizing** layer that organizes previous developments into a cultural artefact, like Euclid's Elements, which may have had various precursors doing part of the work.

A much earlier transition (which may actually be many transitions combined) is from organisms able to acquire and use information about their **immediate** environment including only the immediate past, to organisms that can acquire and store and use information about an **extended, enduring**, structured environment, most of which endures when not sensed, with different locations, routes between locations and contents of those locations including foods of various kinds, shelters of various kinds, conspecifics, and offspring to be guarded, fed, etc.

NB

I use the word "information" in the sense of Jane Austen (and many others before her) rather than in Shannon's sense, as explained in [Sloman\(2013--2018\)](#). Unfortunately Shannon's choice of the word "information" for the concept he was discussing (a form of information capacity, with no reference to information content) has confused many researchers in many disciplines, although I don't think he was confused.

As far as I know, nobody knows how that transition occurred though I have some partial (quarter-baked) ideas about possible intermediate stages. If you know of anything specifically on that sort of transition (a major type of creativity in the mechanisms of evolution of information-processing capabilities) I'd be grateful for references, including things you have written on that.

Requirements for new biological states and processes

It might be thought that the DCKs do not add anything, since everything they produce includes only components that are parts of the FCK. However, this underestimates the importance of new structural relationships and consequential causal powers.

It is certainly the case that anything produced by a DCK, however complex and novel it is, will be physically composed entirely of components of the FCK. This might give the impression that it could be completely described in a language that refers to nothing but the components of the FCK that are its physical parts: the language of fundamental physics (whatever that might turn out to be). But there is a subtle error here, that is very hard to describe, though it can be illustrated by claim that a true sentence describing what is happening in a chess virtual machine running on a computer, for example the sentence "It has detected a threat that is searching for a move that will block the threat" may not be translatable into the language of physics.

Why it is not translatable despite the fact that the virtual machine is *fully implemented* in physical matter is not obvious to most people? Part of the reason is that the original statement about what is happening could be true of the chess machine even if it turns out that our current physical theories are inadequate and therefore the proposed translation does not accurately describe what is going inside the physical machinery.

Closely related to this point is the fact that it is always possible that new physical discoveries or new engineering techniques will allow parts of the physical machinery supporting the chess virtual machine to be replaced with new components that were not mentioned in the original proposed translation.

Even if the alternative implementation techniques were known about and included as a disjunctive component in the description of the machine, the intricacies of the various physical descriptions mentioning newly discovered physical mechanisms or new implementation technologies for computers are irrelevant to the original facts about what was going on in the machine.

This is a bit like fact that the description of drawing as a picture of a circle inside a square remains correct even if the details change in any of a myriad of ways: thickness of lines, colours of various parts of the lines, size and location of the circle, size, location and orientation of the square and so on. The square being in the circle means that no route from outside the circle to a point on the square can avoid crossing the circle boundary. This is a mathematical relationship that can be implemented in (potentially, or actually) infinitely many different ways. But the relationship is invariant to those details.

Similarly the description of the chess programs goals and what it was trying to do to achieve those goals describes a sort of mathematical feature of what is going on that is invariant to an enormous number of possible implementation details. A much more familiar example, which is deceptively simple, is the description of a control mechanism as using a negative feedback control loop, such as a Watt governor, or a thermostatically controlled heater. The difference is that in one case the invariant relationship is a spatial (topological) relationships, whereas the invariant relationship in the chess case involves functional roles and causal powers of sub-processes.

A further complication is that there may be no way for the whole system to be assembled from physical matter without the use of additional materials, including tools, cutting mechanisms, assembling mechanisms, scaffolds, etc. So a full physical description of the finished product will leave out important features of its history that are required for it to exist.

As any architect or engineer knows it is typically impossible to produce an intricate multi-functional machine with complex cooperative parts simply by assembling atoms or molecules in a stock of chemicals. All sorts of new materials (some requiring complex machinery and chemical plants for their manufacture), new sub-assemblies, new externally supported partial constructions are needed, along with sophisticated information-based control systems of various sorts, including living or artificial machines (designers, managers, and workers) that are able to use the latest scientific theories and engineering know-how in proposing goals and finding means to achieve those goals by choosing among options available at various stages of design and construction, and sometimes abandoning one choice in favour of another when things go wrong.

Biological creativity increases spontaneously

It should be clear that we are not talking about *degrees* of creativity measurable on some scale, as discussed in connection with cultural creativity by [Skusa&Bedau \(2002\)](#) but about *structured kinds* of creativity that differ not as bigger and smaller piles of sand do, but as different chemical compounds do, with increasingly complex spatial structures and increasingly complex types of chemical property.

Some researchers have attempted to compare human creativity with creativity in other species, leading to claims such as

Many species engage in acts that could be called creative... However, human creativity is unique in that it has completely transformed the planet we live on. We build skyscrapers, play breathtaking cello sonatas, send ourselves into space, and even decode our own DNA. Given that the anatomy of the human brain is not so different from that of the great apes, what enables us to be so creative? [Gabora & Kaufman \(2010\)](#)

However, long before there were humans or even animals on the planet early life forms transformed the biosphere in spectacular ways that made it possible for animals to exist, for instance producing an oxygen-laden atmosphere as a by-product of photosynthesis, and a climate that made new forms of life possible. Of course the earlier forms of creativity were not *intentional*, but they are still highly complex and profoundly influential.

The kinds of creativity found throughout biology depend on three different types of factor, each of which has many different examples:

1. The existence of a mechanism with rich generative power: i.e. it makes a wide variety of types of structure and types of process possible, without necessarily providing a basis for predicting when each possibility will be realised.

A major feature of science, more important than abilities to predict and to produce falsifiable theories, is the ability to explain types of *possibility*, as explained in [Chapter 2 of Sloman \(1978\)](#).

2. The existence of mechanisms that make certain instances of those possibilities likely to occur, though not all instances need be equally likely.
3. The existence of mechanisms that can select between instances, or types of instance, that do occur -- the function allocated to *natural selection* by Darwin.

That general notion of creativity includes human minds and other animal minds (including, for example, squirrels that defeat "squirrel-proof" bird-feeders, recorded in several online videos). It is a key feature (but an under-rated feature) of developing minds of the young of "altricial" species, including humans, because so much of their cognitive/mental development is *meta-configured*, not *pre-configured*, in the terminology of [Chappell and Sloman \(2007\)](#)

Creativity in Artificial Intelligence

Varieties of computational creativity research

The topic of Artificial Creativity, and more generally attempts to demonstrate creativity in computers, seem to have received increasing attention in recent years and has also generated controversy as indicated by the discussion page associated with the Wikipedia "Computational Creativity" page:

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

https://en.wikipedia.org/wiki/Talk:Computational_creativity

In this document I'll try to explain why "natural creativity", i.e. the forms of creativity found in humans and other animals, are essentially all computational, and will try to identify types of computational creativity that can occur in evolution, in individual organisms, and in past and future AI systems, especially those concerned with explaining, modelling or replicating forms of natural intelligence, in humans and other species.

I shall later try to show why the concepts "creative", "creativity" are examples of parametric polymorphism, discussed in:

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

Roughly what this means is that there need not be anything in common between actions, mechanisms, or products of different kinds that are creative, e.g. creative cooking, dancing, programming, painting, research, composing music, composing novels, engineering, philosophy, etc. What may be common is a collection of relationships to other things (purposes, resources, tools, procedures, possibilities, etc.).

Is it possible for computers to be creative?

It is natural for people to think that because computers are programmed by human beings they can only do exactly what they have been instructed to do, which is no more creative than a human writing a poem or composition by following very detailed pre-constructed instructions, which always produce exactly the same results, or which produce results that vary only because some of the options are selected randomly, as in this random Haiku generator:

<http://www.cs.bham.ac.uk/research/projects/poplog/teach/haiku-demo.txt>

That view about limitations of computers is usually based partly on inadequate theories of creativity and partly on ignorance about the abilities of computers to perform in ways that were neither intended nor foreseen by their authors and which are, to a large extent, responses to unexpected problems encountered while the programs are running rather than behaviours intended or anticipated by programmers.

In particular, many relatively small decisions taken by designers, each with a particular intention (e.g. to meet an engineering need) can, when put together, make new designs possible that none of those designers has previously thought of.

This is a feature of advances in many fields: physics, chemistry, mathematics, engineering, artistic design, architectural design, and others.

Even professional software engineers and computer scientists who know a great deal about computers often fail to understand the full implications of what they know. Since a programming language will typically make infinitely many new programs possible (ignore physical address space limits in computers), computer sciences frequently produce something with far more creative power than they can understand.

That often includes the power to produce new bugs and unwanted behaviours.

There are many different sorts of discrepancies between programmer/designer intentions and what actually happens, some of them comparable to discrepancies between information in an animal's genome at the time of conception and what the animal later does.

In some cases the changes that happen during the running of the system are mainly additions to what was initially specified, but in others the initially specified behaviours are "bootstrapping" behaviours that cause information that happens to be available in the "run-time" environment to replace the initial bootstrapping program. Simple bootstrapping replaces the initial small program with a much larger one, which takes control of the machine. In other cases the program read in may be modified by the 'bootstrapper' to suit hardware details of the machine, local conditions (geographical location, date, time, local language), and other information read in across a local or global network.

Such processes may happen several times, to different portions of a program.

Some of the new run-time information consists of programs written by humans for the specific machine, but a great deal of it may be information picked up from other sources, including automatically stored records of previous events on the same machine, sensor information, information acquired across networks (e.g. incoming email, news channels, weather reports, updates to operating system and other software, and many more).

Some of this post-boot information absorption can be compared with processes like *imprinting* in some species ([https://en.wikipedia.org/wiki/Imprinting_\(psychology\)](https://en.wikipedia.org/wiki/Imprinting_(psychology))), and *language learning* or *absorption of a culture* in humans, although there are still important discrepancies between what computers can do and what intelligent animals can do, some of them discussed below.

Types of post-boot behaviour

To help undermine some assumptions about relationships between an initial program run at boot-time and things that happen later on, I'll present some examples of ways in which behaviour of computers can go beyond intentions of programmers, i.e. ways in which they might be said to be creative. Comparable remarks can be made about the creativity of biological evolution provided by features of the physical universe -- i.e. the "Fundamental construction kit", discussed in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>

Some of these are differences of degree, others differences of kind. Working out implications of complex programs can be very difficult in some cases, and running programs often do more, or less, than their designers intended, not all of it welcome.

1. What programs do can depend on far more than the initial programs/instructions; it can also depend on new input while running that is not under the control of the programmer (e.g. sensory input in the case of robots).

2. Programs can be produced by teams of people who don't necessarily understand all the interactions between their different contributions. The results may be welcome or unwelcome.
3. Programs may be to some extent self-modifying on the basis of problems or opportunities they encounter, and after such modifications can behave in ways that the original programmers did not specify. In some cases that will include creating new names for types of states or processes whose possibility the programmer had not anticipated.

For example, a program that plays a board game may discover that it is in a "lost" situation where it is about to lose and cannot make a move to avoid losing: all the moves available to the program in that situation lead to a lost game, if the opponent recognizes the type of situation.

Detection of such a state could occur after systematically exploring all the currently available move options and finding that all of them lead to lost games. This could cause the program to record the current situation as a "losing state" and all the "forward" moves as "losing" moves in that situation. This detection could trigger a process of working backwards to a previous situation where it had some options that it did not take previously. It may be able to describe that situation as one in which it should never in future make the move that led to the current "losing state" situation.

(This kind of learning was built into a GoMoku program by Elcock and Murray in 1966.)

A program with this and other learning mechanisms could play a lot of games, including games played with itself, and learn a great deal about losing states, and in future play in such a way as to avoid getting itself into such states, while recognising opportunities to tempt its opponent into such states.

The programmers need not know, or be able to predict in any detail what such a program will do after it has played several games, so the program could be said to have a certain kind of creativity insofar as it is able to detect and make use of such states.

However, in the sort of case described, that potential exists because the programmers anticipated that such a mode of learning was possible in that game and provided the program with instructions to enable the learning to occur. We could describe that as "instance creativity": the program creates instances of a type, but it did not create or invent the type, or the instructions required to take advantage of instances.

A further shift from programmer control to program creativity could occur if the programmers designed a more general type of game-playing program, which could be instantiated into many different specific games by providing descriptions of start states, win states, draw states and allowed moves. Then instead of programming it to detect lost states and forcing moves, in a specific game, they could provide it with a generic mechanism for learning such forcing patterns in a class of board games.

Given rules for a new board game in which players place or move tokens on a fixed board, a program would automatically play games according to the new rules and discover for itself which types of board situations are "forcing patterns", i.e. situations in which one of the players cannot win because thereafter the other player can play moves that are guaranteed to lead to a win.

Later the program could be given rules for a game that was unknown to the original programmers, and in that game discover and make use of forcing moves.

This requires a more abstract type of programming. It commonly occurs in software engineering that users write several programs for different purposes, then notice an abstract feature that they have in common and then design a new programming language that simplifies the task of producing new instances of that type of program. An example is moving from writing a parsing program for a specific grammar to writing a parser generator program that can generate a parser when given a grammar, as illustrated in

<http://www.cs.bham.ac.uk/research/projects/poplog/teach/grammar>

A further generalisation might enable a program to learn rules of a new game by trial and error, by being given some initial game components (e.g. a board and some coloured pieces) and a general ability to try placing pieces on the board and accept responses from a teacher, like "legal", "illegal", "win", "lose". Then having worked out the rules it can play the game and start finding patterns that enable it to improve its play.

Such a generic capability would in some ways be like the genetically endowed ability of normal humans to pick up the grammar and semantics of any one of a very wide variety of languages.

4. The previous examples did not include any use of statistics: every situation is legal or not legal, a win for one side, a loss for one side or a draw, and every transition from one situation to another is legal or not legal.

A suitably designed generic program might be programmed with these abstract categories, then by trial and error, using feedback from another player discover the rules of a new game, and then eventually work out how to play well in that game.

A slight generalisation of this sort of capability might enable a program to discover that if it gets into certain situations losing happens often thereafter, and learn to avoid those situations. This is how some programs improve their play by learning to avoid those situations.

However that could be disastrous if some of those situations do not necessarily lead to losing provided exactly the right options are taken later. However, recovery from such an error could occur by playing with a superior opponent who chooses the moves thought by the learner to be bad, then demonstrates a novel path to winning thereafter. In that case the learnt probabilities will be corrected and future opportunities not lost.

That sort of discovery could also come from self-play with a partly random opponent rather than an expert opponent, since a random opponent will occasionally make moves that an expert will make, without knowing why they are good moves, and may then go on to win, through "luck". A good learner can learn from inexpert but lucky players, as well as from experts.

The examples so far depend on the fact that the computer's world is a discrete grid of locations with discrete options, which is true of games like noughts and crosses (tic-tac-toe), draughts (checkers), go-moku, chess, and go. One of the consequences of that feature is that there are many spaces that can, in principle, be checked exhaustively, and that means that computers that are large enough and fast enough can perform optimally. Of course some finite, discrete spaces are too large to be explored by any physical device during the lifetime of a planet, or galaxy, or perhaps even the universe.

There have been many AI projects that are designed to demonstrate some aspect of intelligence in such discrete spaces (e.g. "tile-world" programs, including "Game of life" demonstrations). Although many of these are interesting in themselves, there is no reason to believe that successes demonstrated in such worlds are relevant to explanations of animal intelligence, insofar as animals perform in continuous space-time, or approximately continuous space-time, allowing strings to be tied in knots, with knots pulled tight, shirts to be put on and removed, branches in trees to be woven into resting places, etc.

5. **New layers of virtual machinery**

Viewing complex computing systems as a collection of programs running on a physical computer may ignore the important fact that programs can support a *changing* collection of interacting virtual machines created independently by different designers with different goals, and those VMs may interact in unintended ways with one another and with an external environment connected via input/output devices and networks linked to other computers: what happens within the virtual machines may be in part the result of interactions between many events with different independent sources, and not derivable logically or mathematically from descriptions of physical processes in the machines -- like a national or world wide market.

6. More importantly, instead of a possibly changing collection of externally designed and implemented virtual machines, the original designers may have provided extendable mechanisms and tools able to propose and test new programs and programming tools, with extendable languages for specifying goals and constraints when new programs are designed.

Insofar as the newly designed programs are able to generate new goals and interact in novel ways with the environment some of them may develop new theories about what is happening in the environment and embark on programs to test and enhance the theories. The theories need not be restricted to some initially provided vocabulary: if one of the subsystems records an unexplained collection of observations it might invent new labels for some unknown causes of observed phenomena, and embark on processes of empirical research and theory development that could lead it eventually to a collection of theories about chemistry, electricity, magnetism, electromagnetic radiation, and many more. Some of the systems may start developing theories to explain their own (internal) behaviours and the behaviours of other subsystems with which they interact: inventing notions something like our notions of belief, desire, preference, intention, plan, hypothesis, curiosity, and many more, to label types of internal states and processes that had previously existed, but had not been noticed.

7. The history of human evolution (including cultural evolution) suggests that some of those internally generated subsystems might acquire the kinds of mathematical abilities that in humans eventually led to the development of increasingly sophisticated mathematical theories, conjectures, notations, proofs, and applications to practical problems. (I don't think anyone knows at present whether all of those abilities to discover geometrical and topological concepts, conjectures, proofs and refutations[Ref Lakatos] can be fully implemented in general purpose turing-equivalent digital computers or whether human brains need some other form of information processing as well, possibly based on chemical processes but not implementable on purely digital computers.

For all of that mathematical and scientific enquiry to develop it may be necessary for the pre-existing systems to have values and forms of motivation that are not entirely based on expected positive and negative rewards, in addition to the cognitive apparatus required to represent, ask questions about, theorise about and reason about a wide range of topics. (I think this is obvious as regards mathematical and scientific curiosity, and also some of the motivations in young humans and other animals. For more on this see the discussion of Architecture-based vs Reward-based motivation in [Sloman 2009](#)).

A subset of the sorts of programs that seem to have been produced in humans by evolution might also evolve in the new packages, either in accordance with intentions of designers, or unintentionally. This could provide new reasons for viewing the original set of programs as producing creative processes. A minimal sort of creativity would have to be designed in, e.g. an ability to discover, name, implement, test and perhaps use new theories and procedures.

Already there are programs produced by AI researchers that can be regarded as creative: they may have been designed to perform tasks that require creativity, because they perform tasks that are impossible without creativity. I am not here talking of programs design to perform in an existing performing arts framework, though there are such programs, about most of which I have nothing to say.

For example, as argued below interpreting a picture need not be a simple process of mapping image arrays to descriptions of what has been depicted. Achieving a coherent global interpretation of a complex image may require a great deal of co-operation between different subsystems whose combined operation was not explicitly specified by anyone in advance, anymore than the all the specific goals, plans and actions of a high quality self-trained AI chess program are explicitly specified by anyone in advance.

CLOWES

Some of his ideas about the creativity required in humans and in AI vision programs, can be found in the notes and summaries in this informal biography/bibliography recently appended to his obituary notice: <http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html#bio>

For example a talk he gave in 1971 at the Institute for Creative Arts in London (ICA) and later published, was called "Man the creative machine", and was about the kinds of creativity that computers would need, for example in seeing things and understanding pictures for instance pictures involving entangled couples whose visible body parts are not all connected in the image. (An example is presented here: <http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html#clowes-benthall>)

Around that time he was known for a slogan "Perception is controlled hallucination" though it was not included in any of the publications I have seen. The idea of hallucination here does not refer to arbitrary errors of perception, such as drugs or brain damage might induce, but formation of percepts whose semantic content goes far beyond the content of the sensory data (e.g. images) on which they are based. The "Controlled Hallucination" idea was not entirely new, being close to the idea of von Helmholtz that perception is "unconscious inference" and ideas previously published by Richard Gregory, the vision psychologist (REF), Ernst Gombrich and others.

However, by then concern with creativity was not at all new in Artificial Intelligence, as shown by [Minsky \(1960\)](#).

By the time I published of [Sloman \(1978\)](#), it was clear to me that many sorts of creativity were involved in normal human competences and several chapters indicated some of the problems to be solved, especially Chapter 2 which dealt with kinds of creativity, including ontological creativity, in the development of science, and the creativity required for future machine vision systems.

Several of the earliest researchers in AI realised that varieties of creativity were among the requirements for natural intelligence, and emulations of natural intelligence. Attempts to model creative processes of various kinds were made in the earliest decade of AI, including the Pandemonium system of Oliver Selfridge, and the work on problem solving by Newell, Shaw and Simon, and Saul Amarel's analysis of the 'Missionaries and Cannibals' problem showing that a creative problem solver could discover a new representation of the problem that made a solution much easier to find. In the same decade, Elcock and Murray (then at Aberdeen University I think) showed how a program playing Go-Moku could start discovering "Forcing patterns", e.g. by analysing the state preceding a situation where it lost, and gradually learning more and more complicated "forcing chains".

<http://aitopics.org/sites/default/files/classic/Machine%20Intelligence%201%262/MI1%262-Ch.6-Elcock%26Murray.pdf>

The role of creativity in intelligence may have been unobvious to some commentators at that time, partly because they misunderstood the roles of "heuristics" in improving performance, as Minsky pointed out. Leading AI researchers in the 1960s understood the importance of creativity in avoiding default, potentially explosive search processes, and tried to find appropriate implementations. That included work inspired by Chomsky's observations emphasizing the creativity involved in ordinary language generation and understanding, arising from the essentially unbounded "generativity" of natural language (humans on this view had infinite "competence" but limited "performance").

I was introduced to AI by Max Clowes, a pioneering vision researcher who came to Sussex University in 1969. His work emphasised the need for creativity in ordinary perception (including popularising the slogan "Perception is controlled hallucination" -- related also to ideas of von Helmholtz, the 19th Century polymath, Richard Gregory the psychologist, and Ernst Gombrich, the art historian). The work by Clowes was different from the previously dominant paradigm of 2-D pattern classification, a very different task from constructing 3-D interpretations of 2-D images. The latter was an inherently creative process. One of his important papers was based on a lecture he gave to the ICA in 1971, making this point:

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

There is more information about his work in this tribute (he died in 1981):

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61>

In view of all this I was completely taken aback when I recently encountered a suggestion that the study of computational creativity was something new.

Eventually I think I worked out what had happened. There is something new, but it is only one rather narrow approach, among a variety of ways of relating computation and creativity. As a result, I have concluded that it is useful to distinguish computational research on creativity into the

following categories (using tentatively suggested labels):

Type 1. Intrinsic creativity

This includes aspects of creativity involved in ordinary perception, learning, planning, plan execution, speaking, understanding, making things, designing and making artefacts and tools, doing scientific research, etc. The need for creativity comes from task features requiring combinations of different kinds or sources of information, also from combinatorial complexities and the need to handle noise, and, in the case of AI vision systems, incomplete information caused by occlusion (including self-occlusion), distance, poor lighting, looking through shrubbery, perception of things in motion, perceiving while the perceiver is in motion, and other sources of difficulty intrinsic to intelligent performance. This category could also be called "natural creativity", or "embedded creativity" because creative processes are embedded in many cognitive functions. (Explaining the sense in which they are all creative is a topic for another occasion.) There are now many online videos of squirrels, nest-building birds, human toddlers, and others exemplifying this kind of creativity.

The need to implement various kinds of intrinsic creativity is a constant source of challenges for AI system designers. Studying and modelling the mechanisms is a challenge for scientists, including AI scientists. Replicating them is a challenge for AI engineers, since robots and other intelligent machines will constantly have to exercise this sort of creativity. The slogan "perception is controlled hallucination", made popular by Clowes in the 1970s summarises a recurring type of challenge for AI researchers in vision. (I have learnt from Keith Oatley that the key idea had been pointed out earlier by Hippolyte Taine in 1882). I'll return to the need for creativity in vision later, because it concerns a serious reporting error.

Type 2. Expressive creativity, perhaps also describable as "talent creativity"

The second sort of creativity is more publicly visible than the first kind: it includes aspects of creativity regarded by humans as special because they need special competences or gifts or talents, and those needs do not frequently emerge from standard processes of daily life and survival. These would include specialised, and not necessarily widely shared, musical, poetic, artistic, mathematical, technical or scientific creativity, including abilities to invent stories, design new machines, create scientific explanations or hypotheses or questions, etc. AI research on this sort of expressive creativity includes the work on painting by Harold Cohen, the work on musical composition by David Cope, and more recent work by many people trying to get computers to compose music, make pictures, compose poems, create puns, etc. I do not claim that there is a sharp division between this and the previous type of creativity. It may be worth further sub-dividing both categories according to types of biological function and their roles in evolution.

Type 3. Standardised test-based creativity

The third kind of creativity, seems to be a recent interest of psychologists who have devised collections of tests, so that individuals can be given creativity scores, which may be thought to be related to assessments of creativity of the first two kinds, though I am not sure about that.

In principle such tests could include tests for expertise in mending broken machines, or in composing or completing interesting poems or pictures. (I don't know whether the "completing"

tasks have actually been used for such tests, though I could imagine someone using them.) Computational work on this kind of creativity depends crucially on the prior existence of standard tests for creativity that are capable of assigning scores to those who are tested, because there are right and wrong (or perhaps good and bad) answers.

(Boden's distinction between H creativity and P creativity can be applied to all three categories, though there will be less scope for H creativity in Type 3 tests.)

Recently, it seems that the third kind of creativity has been referred to as "computational creativity", which is potentially confusing because the other two varieties are also clearly computational: they involve information processing capabilities in humans, and computers are being used to model or emulate them.

The need for Type 3 creativity typically does not arise intrinsically in purposive thinking or action by an individual, as in Type 1. Individuals encounter it mainly when they are being tested for what psychologists regard as creativity, though some of the tests may be inspired by examples that were originally of Type 1.

Insofar as the need for Type 3 creativity does arise in real life there does not seem to be anything systematic about it: different individuals stumble across and solve different problems, some of them of great practical importance, others of minor importance, or in some cases as sources of entertainment. I think the problems of intrinsic creativity are much deeper, more systematic and will turn out to be far more important for understanding how minds work, how they evolved, how they develop, and how they are implemented in brains.

Work on Type 3 creativity tends to be associated with evaluation processes that produce a numerical measure of creativity. But this requirement tends to be based on a narrow view of science: we can classify all sorts of things by producing structural descriptions (e.g. molecules) or combinations of structural and functional descriptions (e.g. species) without assuming that all the categories are distinguished by numbers.

What I have called Type 1 (intrinsic) creativity need not be measurable in any non-arbitrary way, any more than varieties of plants are measurable. There just are many kinds of plants. Likewise there are many kinds of Type 1 (intrinsic) creativity, embedded in natural information processing.

The same can be said of creativity of Type 2 (Talent creativity). There are certainly partial orderings: some musical compositions are better than others, but I hope it would be agreed that an attempt to place the compositions of all the great composers on a scale of creativity would be absurd, and trying to justify a numerical measure comparing the creativity in a great poem with the creativity in a great string quartet would also be absurd.

This does not rule out *partial* orderings. Some examples of one category may be rated more valuable than examples of another category without presupposing any *measure* of value applicable to everything in both categories.

Many years ago, after studying reports in *Which?*, the UK consumer magazine, I proposed an analysis of "Better" as a concept with *parametric polymorphism*, though I did not know that label at the time. I think something similar could be said about the concept of creativity, but I'll not explore that till a later date.

In view of my arguments about the inapplicability of measurements and total orderings to examples of Type 1 and Type 2 creativity, it follows that insofar as Type 3 creativity is restricted to what psychologists (or others) can measure in some way, the work on Type 3 creativity will not necessarily address the kinds of deep, pervasive, and systematic creativity that make up human and animal intelligence.

A focus on *measurable* creativity seems irrelevant to the kinds of creativity that have been recognized as important for AI as science: there are many forms of intrinsic creativity but no reason to believe that they are all assessable on some scale of creativity. So the emphasis on *measurement* restricts the scope of the research. There is nothing wrong with restricting a research project in that way, but excludes much interesting research on kinds of creativity that cannot be measured as suggested.

[Insert something about creativity in visual perception, including interpretation of diagrams.]

The most visible type of AI creativity research

Type 2, *expressive/talent creativity* is what is perhaps best known as a type of creativity studied in AI, by people trying to design AI systems to produce novel pictures, music, puns, stories, mathematical discoveries, etc., or musical or other performances. In some ways this is the most "visible" type of creativity studied in AI.

I have been aware of AI research on Type 2 creativity, but it has never been one of my main interests since I am more interested in the kinds of everyday "intrinsic" (Type 1) creativity displayed by humans and other animals as part of their natural activities. However, I have found some of the examples interesting, e.g. Harold Cohen's painting program AARON, described briefly in Section 2.3 of the paper, because I think that that work actually addresses some of the requirements for Type 1 creativity, though I am not sure. A few other examples of "talent creativity" (Type 2) are mentioned in that section, but without any analysis or critical comment of the sort I was expecting. Those are the types that tend to attract most attention from the general public.

A new field of "computational creativity" research

Creativity and AI

From the very earliest days of AI the leaders of the field recognized the role of creativity in problem solving (not least because they were aware of prior work on creativity, e.g. by Kohler, Wertheimer and others), but also because it was already obvious when Minsky wrote [Steps \(1960\)](#) that various forms of creativity (including learned creativity) were required, e.g. in order to defeat combinatorial explosions and other problems.

This recognition of the need for creativity in AI systems was linked to ideas about Type 1 creativity of human minds in everyday life stressed also in contemporary work in Linguistics by Chomsky and his fellow workers who (rightly) emphasised the creativity involved in most ordinary language generation and understanding, arising from the "generativity" of natural language.

The AI-related interest in Type 1 creativity draws attention to important forms of creativity hidden in many human and non-human achievements in their everyday life -- e.g. a crow working out where to put the next twig in a part-built nest, or a pre-verbal toddler performing a topological experiment with a pencil and a sheet of card containing a hole, as shown in a video presented here:

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

As Chomsky implied, most of the sentences in a document like this will be novel creations, including, I suspect, this one, for which google found "No results". The creativity involved in producing such a sentence is matched by the creativity involved in comprehending it "on the fly", which I expect is achieved by everyone who reads this report, though I doubt that there currently exists an AI program that could read and understand such a document (because difficulties pointed out by Bar Hillel around 1960 have not yet been overcome by natural language researchers).

Moreover, as explained above, my own introduction to AI over 45 years ago, by Max Clowes, emphasised the importance of creativity in natural intelligence, and especially creativity in visual intelligence. My interest in Type 1 intelligence also emerged from my interest in modelling the ability of humans to make discoveries in geometry and topology, as a way of supporting Kant's philosophy of mathematics, which is what led to the 1971 paper making the distinction between Fregean and Analogical representations, mentioned in the paper.

Is a new field of "Computational Creativity" a diversion?

It is often the case that researchers like to label themselves as doing something new, whose importance was not recognized in so-called "classical" research. In my experience that is often a form of self deception in the pioneers of the new work (e.g. in fashions for embodiment, enactivism, situated cognition, ...). And it can divert attention from hard, important, problems.

The focus on laboratory tests for creativity and emphasis on *measurement* of performance rather than *explanatory analysis* of varieties of creativity, seems to me to lead to impoverished psychological research, compared with, for example, identifying the variety of types of creativity required at various stages of development of very young children, including their linguistic development -- which is often interpreted as passive learning, but is clearly often highly creative, as shown by the deaf children in Nicaragua who invented their own rich sign language because their teacher could not meet their needs, and twins who invent their own language. See <https://www.youtube.com/watch?v=pjti0FuNf8>

The importance of creativity in young children has been noticed by several educationalists and developmental psychologists, e.g. John Holt, Piaget, Vygotsky and others. However, the challenge of modelling such creativity is beyond the current state of AI (for various reasons). I hope that will one day change.

The need to know about the history of AI

For about half a century AI researchers have been struggling with ill-structured problems. For example, such problems are unavoidable in AI research in natural language, in vision, and in getting robots to act in natural environments.

However for most of that time the computational resources available to AI researchers in their laboratories have been miniscule compared with the computational power now available in everyday devices, including mobile phones, tablets, etc. E.g. Sussman's HACKER program that discovered new ways of anticipating and avoiding programming errors could not all fit into the AI computer at MIT in the early 1970s, so it could not all be tested at once. It may be difficult for young researchers now to understand how limiting those computational resources were for "classical" AI. When I visited Edinburgh in 1972-3 there was great rejoicing at the doubling in size of memory in

the departmental computer, from 64KBytes to 128Kbytes.

Despite those limitations the Edinburgh robot Freddy (using an additional small computer for handling I/O) was able to deal with a cluttered pile of wooden parts and assemble a toy car or boat. Sometimes this required it to grab a portion sticking out of a pile and move it from side to side, to separate parts that previously could not be seen clearly. In that case the creativity had to be installed by a programmer (Chris Brown), but it was already clear that AI robotic problems in "real life" would require considerable creativity of that sort and other sorts. There are videos available of Freddy:

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

In other areas, e.g. the problems of AI vision systems as discussed by Max Clowes and others in the early 1970s and 1980s, it was very difficult even to get agreement on the goals of the research, with some researchers regarding vision as largely a process of pattern recognition, others (e.g. Marr) emphasising derivation of 3-D structural descriptions from (often cluttered and noisy) images, and yet others, partly inspired by Gibson regarding the functions of AI vision as identifying affordances. My own work at that time included the role of vision in mathematical discovery and reasoning, a problem as ill-structured as any.

Is Type 3 (test-related) creativity over-rated?

To be continued.

Note:

I have learnt much, over many years (since around 1965), from Margaret Boden, the leading philosophical thinker about creativity and its relations to AI.

<http://www.ruskin.tv/maggiieb/publications.asp>

I am also grateful to Ana-Maria Olteteanu whose PhD thesis, focused mainly on Type 3 creativity, indirectly triggered the lines of thought being developed here.

REFERENCES AND LINKS

M. A. Boden (1990/2004) *The Creative Mind: Myths and Mechanisms*, London, Weidenfeld & Nicolson, 1990, 2nd Edition, Routledge, 2004

Jackie Chappell and Aaron Sloman

Natural and artificial meta-configured altricial information-processing systems
International Journal of Unconventional Computing, 3, 3, 2007, pp. 211--239,

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717>

See [the diagram](#) on the main M-M web page.

M.B. Clowes, (1973), Man the creative machine: A perspective from Artificial Intelligence research, *The Limits of Human Nature*, Ed. J. Benthall, Allen Lane, London, (Based on a lecture given to the ICA in 1971).

Simon Colton, 2016, 'Computational Creativity', Talk at the *Rustat Conference on Superintelligence and Humanity*, June, 2, 2016,

<https://metamakersinstitute.com/2016/06/06/rustat-talk/>

<https://www.jesus.cam.ac.uk/articles/machine-superintelligence-and-humanity-rustat-conference-report>

David Deutsch, 2011, *The Beginning of Infinity: Explanations That Transform the World* Allen Lane & Penguin Books,

Euclid, *Elements*

<http://aleph0.clarku.edu/~djoyce/java/elements/toc.html>

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

Gabora, L. & Kaufman, S. (2010). Evolutionary perspectives on creativity. In *The Cambridge Handbook of Creativity* Eds J. Kaufman & R. Sternberg, pp. 279-300, Cambridge UK: CUP

<http://arxiv.org/pdf/1106.3386.pdf>

Tibor Ganti, *The Principles of Life*, Eds. Eors Szathmary & James Griesemer, OUP, New York, 2003, Translation of the 1971 Hungarian edition, (Review by Korthof:

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

Stuart A. Kauffman (1993) *The Origins of Order: Self-Organization and Selection in Evolution* Oxford University Press

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

Giuseppe Longo, Mael Montevil and Stuart Kauffman, (2012) No Entailing Laws, but Enablement in the Evolution of the Biosphere, *Proceedings of the 14th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO '12*, ACM, New York, NY, USA, pp. 1379--1392,

<http://arxiv.org/abs/1201.2069v1>

Marvin Minsky (1960), "Steps toward artificial intelligence",

<http://web.media.mit.edu/~minsky/papers/steps.html>

Reprinted in *Computers and Thought*, Eds. E.A. Feigenbaum and J. Feldman, McGraw-Hill, pp. 406--450, 1963

David Mumford, Grammar isn't merely part of language, Oct, 2016, *Online Blog*,

<http://www.dam.brown.edu/people/mumford/blog/2016/grammar.html>

Erwin Schrödinger, 1944, *What is life?*, CUP, Cambridge.

I have a collection of extracts relevant to this discussion, with comments, here (still incomplete in March 2016).

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

Andre Skusa and Mark A. Bedau (2002) Towards a Comparison of Evolutionary Creativity in Biological and Cultural In R. Standish, H. Abbass, and M. Bedau, eds., *Artificial Life VIII*, MIT Press.

<http://www.reed.edu/~mab/publications/papers/bedauALife8.pdf>

Arthur L. Samuel (1959) Some Studies in Machine Learning Using the Game of Checkers in *IBM Journal*, Vol. 3, No.3. July, 1959.

<https://www.cs.virginia.edu/~evans/greatworks/samuel1959.pdf>

Aaron Sloman (1978), *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind*, Harvester Press (and Humanities Press), Hassocks, Sussex. A new online edition is freely available here:

<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp>

A. Sloman, (1984-2014) Experiencing Computation: A Tribute to Max Clowes, in *New horizons in educational computing*,

Ellis Horwood Series In Artificial Intelligence, Ed. Masoud Yazdani, pp. 207--219, Chichester, Expanded with bibliography and biography 2014:

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61>

A. Sloman, 2009, Architecture-Based Motivation vs Reward-Based Motivation, in *Newsletter on Philosophy and Computers*, American Philosophical Association, 09, 1, pp. 10--13, Newark, DE,

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

Aaron Sloman, 2013--2018, Jane Austen's concept of information (Not Claude Shannon's) Online technical report, University of Birmingham,

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

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

Aaron Sloman (2017-...); Construction kits for evolving life: The scientific/metaphysical explanatory role of construction kits.

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

(This extends a long paper on Construction Kits published in a Springer collection in 2017), in *The Incomputable: Journeys Beyond the Turing Barrier*, Eds. S. Barry Cooper and Mariya I. Soskova,

<http://www.springer.com/gb/book/9783319436678>

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