WORK IN PROGRESS: LIKELY TO CHANGE

The Meta-Configured Genome
(Multi-layered, multi-stage, parametrised, epigenesis)

Aaron Sloman
http://www.cs.bham.ac.uk/~axs/
School of Computer Science, University of Birmingham

Based on collaboration with
Jackie Chappell
https://www.birmingham.ac.uk/staff/profiles/biosciences/chappell-jackie.aspx

14 Oct 2019
A new alternative (richer?) introduction to this topic
Available at
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/meta-config/
Including a new short (4.5 Min) video explaining the main diagram.
Some of the features of DNA, genome duplication and gene expression that
contribute to the Meta-Configured Genome are summarised in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/dna-uses.html (also **PDF**).

Installed: November 2017
Last updated: 14 Oct 2019
This paper is available in html and pdf formats:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.pdf

This is part of the Turing-inspired Meta-Morphogenesis project
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html
which, in turn, is a recent (2011) development the CogAff (Cognition and Affect) project begun three decades earlier:
http://www.cs.bham.ac.uk/research/projects/cogaff/
A partial (possibly out of date) index of discussion notes in this directory is in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html

SUMMARY
Biological evolution produces many patterns of evolutionary change and individual development, including changes in size, shape, strength, development of body parts, and changes in information processing -- what information is acquired, how it is manipulated, whether and how it is stored, what it is used for, how uses are controlled or improved, whether it is communicated, and many more.

Some evolved mechanisms or competences start working at a very early stage of development and continue throughout life, with minor changes, e.g. construction of a heart and control of heart-beats. Other mechanisms may be modified by learning processes, and other control processes

responding to aspects of the environment during development, e.g. changes occur in control of sucking, swallowing and later chewing, and (more obviously) in control of limb movements, attention direction (e.g. using eye gaze), and how acquired resources, including information, are used.

Differences that emerge may be influenced by changing body size, shape, and muscular strength, and also by aspects of the physical and social environment.

Many of the evolved mechanisms are directly concerned with production and controlled use of body parts, including sensing and motion control mechanisms. Others are more abstract, concerned with aspects of information that can have multiple uses and are not therefore permanently associated with particular sensor and motor subsystems but may become engaged when needed. E.g. route finding is relevant to hunting, scavenging, fetching nesting or shelter building materials, seeking a mate, etc. So some forms of information need to be permanently connected with particular body parts and their uses, whereas others are used as resources to be selected when needed, perhaps using resource catalogues and indexing mechanisms.

Gene expression must account for aspects of: production and growth of the physical body, production and development of mechanisms for acquiring, manipulating, storing and using information of many kinds, including combinations of abilities to control behaviour, e.g. by generating new motives, and abilities to acquire, use and store sensory information, and abilities to produce changes in how information is processed, stored and used.

During their lifetime organisms undergo many physical changes, including spectacular physical changes from a fertilised seed to a recognizable member of a species; and changes in behaviour from early reproductive processes of growth and development largely determined by the genome and later on changes in behaviour increasingly controlled by the organism's needs and the resources, opportunities and threats in the environment, over increasingly extended times and spaces, and in some cases control of conspecifics and other organisms.

An almost helpless human baby able to do little more than absorb nutrients and change size and shape as internal organs develop, may later control the behaviours of large numbers of humans, machines and physical materials including collaborating in hunting for food, fighting human competitors or intruders, and later on constructing large and complicated machines, buildings, and teams of workers acting for extended periods of time over extended spatial regions, e.g. constructing new types of machinery, much larger buildings, extended transport mechanisms and networks, thereby producing massive changes in the physical environment. Others may collect information, process it internally over an extended period, and in some cases produce powerful theories, create influential works of art, develop new technologies and new ways of thinking and alter the minds of others, without directly causing any notable changes in their own physical environments.

Biologists (especially ethologists), psychologists, linguists, educationalists and others have collected information about such developmental changes and tried to produce explanatory theories about why they occur and the mechanisms that make them possible. In parallel with all that rulers, generals, slave-owners, employers and some visionaries have found ways to harness and steer such capabilities towards their own ends, or supposed greater ends.

Over many generations the accumulated effects of all these individual processes can alter large portions of the surface of planet beyond recognition, in some cases destructively. They also alter the individual behaviours of humans, including the food they eat, the clothes they wear, the buildings they create, the activities in which they engage, the languages they use and the ways they interact with within and between groups of various sizes.

How can a species with collections of genes that have been largely unchanged over thousands of years produce such complex changes in themselves and their environments, unlike other species, including intelligent species such as apes, nest-building birds, and hunting mammals.

Part of the answer is that human intelligence is not something fixed by the genome, though it is enabled by the genome, which, in turn, is produced by biological evolution, over millions of years. But there is something about the human genome that made possible very rapid changes in abilities and behaviours without changing the genes. Clearly this depended on newly acquired information being transmitted across generations without altering genes. How?

The obvious answer is that the development of human communication abilities made it possible for newly acquired factual and procedural information to be transmitted across generations without altering genetic mechanisms. This was a complex process that involved generations passing on not only verbally expressed information, but also designs for tools, machines, clothing, buildings, weapons, transport mechanisms, food production methods, etc. Crucially, the machines include machines for making new more complex machines.

CONTENTS LIST (to be extended)

- Evolution as epigenesis
- Other species
- Disembodiment of cognition evolves
  (Epigenesis of evolutionary mechanisms}
- Types of competence, rather than types of species
- REFERENCES AND LINKS

Overview/Summary

Alan Turing's work inspired the meta-morphogenesis project, of which this document is a part: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html (or pdf), But I don't agree with everything he wrote.

In particular in Turing(1950), he wrote (in Section 7 *Learning Machines*) *"Presumably the child-brain is something like a note-book as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets. ... Our hope is that there is so little mechanism in the child-brain that something like it can be easily programmed."*

I'll try to show that that claim is mistaken: it grossly under-estimates the achievements of biological evolution in producing aspects of the human genome (and genomes for other more or less intelligent species) concerned with multi-layered production of information-processing mechanisms required for many of the capabilities of human and non-human minds.

A variant of Turing's conjecture, without the "rather little mechanism" claim is the Modular Mind theory, which proposes that human cognition (or more generally, animal cognition) uses a collection of information processing modules specified in the genome, available at an early stage (before or soon after birth) to drive processes of learning, development, and control of behaviours. Partial surveys of "modularity" theories can be found in Robbins (2017) and https://en.wikipedia.org/wiki/Modularity_of_mind. Examples of such theories were presented by Fodor (1975) and (1983), and others.

Annette Karmiloff-Smith launched an attack on such modular mind theories, e.g. in (1992) and (2006), arguing that competences that were thought (by some researchers) to depend on innately specified modules emerge from complex interactions between neural mechanisms during individual development. Abnormalities in such developmental processes can produce abnormalities in the resulting competences.

Another alternative, different from all of (a) Turing's "simple mechanism" hypothesis, (b) theories about collections of innate modules, active from birth, and (c) the fortuitously emergent competences hypothesis, is what I have been referred to as the *Meta-Configured Genome* hypothesis schematically summarised in Figure EPI below.

The core claim is that the genome includes specifications for "delayed" modules of varying complexity that (a) do not become active until some time after birth, when information has been acquired and stored as a result of activities of earlier modules providing information-gathering competences, (b) make use of the previously acquired information to fill gaps in the specifications of the delayed modules. In other words, the genome does not fully specify the delayed modules, leaving gaps to be filled partly by information acquired before the modules become active and and

partly by ongoing interaction with the environment (in collaboration with other modules).

These late-activated, incompletely specified modules (the meta-configured modules will have different effects in different environments, because the previously developed competences will have acquired different information from the environment, used to act as parameters (gap fillers) in the late developing competences.

The best-known, most obvious (and perhaps most spectacular) examples of such mechanisms are observed in the diversity of ways in which relatively late developing linguistic competences (e.g. abilities to use relatively complex syntactic constructions) make use of previously acquired linguistic competences at various levels, including the basic language sounds, elementary syntactic forms, early vocabularies, and early semantic contents, so that the same communicative action is achieved in very different ways. For example (thanks to Google translate), an English speaking child's utterance of the words "I fell over on my way to school this morning" could perform the same communicative action as a German child saying "Ich bin heute Morgen auf dem Weg zur Schule hingefallen", or an Italian child saying "Sono andato a scuola stamattina." where not only are the words used different, but the corresponding words are expressed in different ways, and in some cases different numbers of words are required to say the same thing in different languages.

> (Compare the claim attributed to Bar Hillel in the 1960s that a computer-based translation program asked to translate "The spirit is willing but the flesh is weak" into Russian and back to English might produce "The whisky is fine, but the meat's not so good".)

> Human language development has many layers, including late-developing syntactic, semantic and pragmatic linguistic competences, whose details will be partly specified by the genome in language-independent ways, and partly based on information acquired both from the use of previously developed competences and also from interaction with the physical, social and linguistic environment as the child develops. An automatic translation system whose design bypasses those stages is likely to be seriously flawed.

The "Meta-configured genome" idea was developed in collaboration with Jackie Chappell after she came to Birmingham from the Ecology laboratory in Oxford, in 2004. Our first paper on this topic was Sloman and Chappell (2005a) followed by this invited journal paper, Chappell and Sloman (2007). The ideas have been modified and extended from time to time since then.

This theory claims neither that all information about cognitive mechanisms is specified in the genome and available from birth, or earlier, nor that the information is acquired by powerful learning mechanisms available from birth and used throughout life (as in Turing's speculation quoted above), nor that the capabilities that arise during development, rather than being genetically pre-specified, emerge from interactions between developing/growing pre-specified neural subsystems, as suggested by Annette Karmiloff-Smith in several publications referenced below.

Instead, the meta-configured genome theory claims that some genetically specified late developing modules are incompletely specified, since they have gaps (i.e. parameters) that are filled by information acquired during early development and later, and that the activation of these modules is explicitly delayed until after earlier modules have had time to acquire information that can be used to provide parameters for the later modules.

Instead, we propose that different kinds of information, including control information, and information about types of information and types of use of information, at different levels of abstraction (i.e. various kinds of meta-information, with information gaps), acquired by the species during different evolutionary stages in its history, are made available to the developing organism at different stages of development, when the inherited abstractions can be combined with detailed fillers, or parameters, acquired at different stages of development.

We call such a genome "a meta-configured genome" (MCG). Something like this theory is required to explain how a common genetic heritage can account for humans acquiring very different languages at different stages of social history and in different geographical regions. However, the idea is not restricted to language or to humans: our claim is that many aspects of intelligence, in humans and other species, make use of MCG mechanisms, though the products in humans are most spectacular, partly because of the evolution of self-reflective (meta-cognitive) mechanisms that allow various kinds of internal information processing to be detected, recorded, reflected on and later modified -- one of the themes of the CogAff (Cognition and Affect) project mentioned above.

In some cases when such abstract genetic information is expressed, it needs to be combined with previously acquired information that can vary across cultures and geographic locations. This information is acquired at earlier stages of development, in some cases on the basis of what I call "architecture-based motivation" (ABM), in contrast with "reward-based motivation". ABM mechanisms include reflex behaviours whose main function is to acquire information that can be stored in case it is useful later, even though the individual has no idea that this is the source of apparently pointless motives (as can often be observed in children). So the meta-configured genome mechanisms depend crucially on the architecture-based motivation mechanisms that allow contents of perception to trigger behaviours that are not motivated by any anticipated reward, since the learner cannot have any idea why these information-gathering actions will produce results that are crucially important at later stages of development. This use of architecture-based, rather than reward-based, motivation is explained more fully in Sloman(2009-2019).

As far as I can tell these ideas contradict all widely discussed theories of learning and motivation.

Moreover, I suspect that if Turing had worked out in more detail, his ideas about mathematical discovery depending on mathematical intuition as well as mathematical ingenuity, this might have drawn his attention to complexities he apparently did not notice. I have tried to summarise his views on that in:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html (also pdf).

Partial summary so far: Most philosophies of learning and science assume there's a fixed world to be understood by a learner or a community of scientists, or other products of biological evolution. However, evolution, in collaboration with its products, is metaphysically creative, and constantly extends the complexity of the world to be understood and constantly extends the types of learners with new powers. Its more advanced (recently evolved) learners have multi-stage genomes that produce new powers within a learner, suited to extending the powers produced by the earlier stages and their environments, partly mimicking - in hugely accelerated form, the earlier discoveries of evolution, partly extending them to take account of discoveries in earlier stages of development. Jackie Chappell and I call this the Meta-configured genome.

Human language as an example

The genetic basis of human linguistic competence is a striking example of this theory. It seems that any newborn human infant is capable of learning any one (and in many individuals more than one) of several thousand different human languages, including spoken languages, sign languages, and written languages, as well as recently developed specialised "technical" languages for mathematics, science, engineering, and computer programming.

In spoken languages there is rich diversity at different levels of complexity in the first few years of life of a young human, where the types of pronunciation (accents), vocabularies, syntactic (grammatical) structures, and semantic contents expressible vary enormously across human cultures, and across individuals within a culture.

If Turing had known all that, or if he had paused to reflect on knowledge that he already had, I don't believe he would have suggested that something like a child brain could easily be programmed. Perhaps he can be excused because he wrote before the work of Chomsky and others regarding the computational features of natural languages. I also suspect that he had not spent much time watching the development of pre-verbal children.

Designing designs

[This section, related to the theory of evolved construction kits Sloman[2017] is still under development.]
The designer of a programming language, or sophisticated software production package, cannot know about, and does not need to know about, all the applications for which the programming language will be used. That is because such languages and packages allow parameters of various kinds to be specified by developers and the package enables those parameters to be combined to generate a new useful structure or package, which may itself have parameters that differ across applications.

The parameters required by the original system or its products may be simple entities such as numbers, labels (strings of characters), bits set on or off, or complex entities such as sets of numbers, other packages to be combined, or in some cases interfaces to hardware devices, such as cameras, microphones, keyboards, screens, or links to various parts of a complex functioning entity controlled by the new package, such as a chemical plant, an airliner automatic landing system, an assembly robot in a factory, a home security system, etc.

Biological evolution also produces designs for abstract re-usable structures or mechanisms that can be used in the production and functioning of individual organisms, or groups of organisms, as a result of insertion of parameters obtained either at later stages of evolution, or during development of individual organisms, or through formation of colonies, cultures or societies in which individuals interact, including transmitting information.

In these cases also, some parameters used for control decisions during development or during particular actions may be numerical values or scalar measures, e.g. temperature, concentrations of certain chemicals, sizes or weights of body parts, or speed of linear motion or angular rotation.

But many parameters in a structured plan or template will themselves be instantiated by structures with parts and relationships, for instance a structured plan for action, may include some labels for unspecified steps, which can be replaced, during execution, with more or less complex actions or objects with previously unforeseen details, e.g. buying a ticket, making a new object, repairing

something damaged, making a new sub-plan, collaborating or competing with conspecifics, etc.

The new component selected or created during execution of a plan may be a physical object, a route, an internal information structure representing a perceived physical structure or process, a theory about how something in the environment works (e.g. an unfamiliar window catch), a structured question to be answered through active perception, experiment or exploration, or an intention to be achieved by performing an action or carrying out a plan involving multiple sub-actions.

These ideas are applicable not only to plans consciously created and used by intelligent agents, but also to processes of development of organisms partly under control of a genome. For example, during physical development, gene expression mechanisms may use structured "inputs", in the form of information about the organism's physical environment, social patterns, or molecular structures, obtained from the organism's environment (or the mother) e.g. through food ingested by the individual or from contents of an egg developing in a shell, or via a mother's womb or milk supplied after birth, etc.

Some parameters are simple "atomic" entities, such as measures or labels: their internal structure does not matter, only their presence or absence, or location in some ordering or set of alternatives. Other parameters are complex, including complex molecules whose parts and relationships play important roles in their use.

Many different structures can be found in such complex parameters inserted "at run time" into a plan. For example, the inserted parameters may be linear chains of linked "atomic" items, such as a complex symbol made by concatenating simple symbols. Or a complex entity inserted into a plan may be a "tree" structure, where there is a "root" entity and every other component is reachable from the root via a unique path through a branching network of nodes and links. In more complex cases a plan item may be a "molecular" network structure with parts that are related to other parts that are simultaneously related to several other parts, with loops in the chains of relationships. Such non-linear, non-tree-like structures often occur in complex chemical structures, indicated by diagrams with loops, and also in maps depicting a network of routes.

There may be several distinct routes between two nodes in such a network structure (referred to as a "graph" by mathematicians, but not in the same sense as a "graph" showing relationships between inputs and outputs of functions).

If evolution, like human designers of software packages, is able to derive generic re-usable abstractions from specific solutions to problems, then more abstract products of evolution can be instantiated (e.g. by setting parameters) for use in contexts in which they did not evolve. That process of instantiation of a complex abstract structure by insertion of other complex structures, is not like a minor variation in size, weight, colour, speed, etc. If it occurs during development of individuals it can produce individuals with novel features that are not minor variants of other individuals sharing the same genome. Complex structural variations in languages acquired in different cultures illustrate this.

If similar variability occurs in a reproductive process that constructs new genomes it can produce new classes of individuals that differ significantly from others with the same ancestry.

This illustrates biological evolution's powerful use of "compositionality" discussed further in Sloman(2018,compositionality).

The requirements for such processes are complex and subtle but it is clear that evolution found such powerful abstractions that are re-usable in new contexts long before human engineers did.

In many cases these developments included "discovered" mathematical generalities that are encoded in a relatively abstract form in genomes, able to be instantiated with different parameters both during development of individuals and when varied across species.

E.g. during individual growth (changing size, shape, strength etc.) and learning (acquiring/creating new information structures at various levels of abstraction) individual development mechanisms can combine information (parameters) from the environment with generic information in the genome to produce instances of a species tailored to the specific environment.

The need for structural variation within and across instances of a genome

The longer an individual lives the more likely it is that the environment will change during its lifetime either because of local changes or because of migration to new locations, or due to other factors such as learning by prey or predators, or rapid evolution of other organisms (e.g. bacteria evolving resistance to medication).

Environmental pressures and opportunities on a species can be changed both during lifetimes of individuals, and more gradually across generations, by climate changes, geological changes (e.g. earthquakes, volcanic eruptions, sea-level changes, floods, etc.), invasion or migration of individuals of other species, removing resources and/or creating new threats and opportunities for previous inhabitants, diseases, crop failures, etc.

All of these cases provide opportunities for benefiting from adaptability both in genomes and in individuals. Resulting features of gene-expression mechanisms can lead to important new features of the interplay between evolution of new species, or new modifications of existing species, and advantages of new mechanisms of epigenesis: genetic influences on individual development.

Some new threats are so drastic that it is impossible for a genome to prepare for them, e.g. impact by a huge asteroid or huge fluctuations in cosmic radiation. Others may fit into a space of possibilities that allows defensive measures to be taken, or new opportunities to be grasped, by alteration of parameters derived from the new situation, in generic mechanisms.

If the parameters can be systematically derived from perceived structures in the new situation, such a process can be much faster than statistical search for appropriate parameters as happens in many current AI learning mechanisms.

A toy example: if you can map the perceived width of a new object into an appropriate extent of jaw opening then you don't need to spend a long time trying out different widths of opening.

If individual members of a species have enough previously evolved generic design features into which parameters derived from perception or a small amount of interaction with new objects or situations can be inserted, then this can significantly reduce the need for extended interactions searching for appropriate parameters: it can dramatically speed up developmental or learning processes. Discovering appropriate parameters for use in an available abstract structure can

substantially reduce the need for "mindless" use of statistical regularities -- blind empirical learning.

Warning Note
This (Kant-inspired) analysis leads to a problem for scientific research on learning and development in humans and other animals. Statistical correlations discovered in observations and experiments may be relevant only to a narrow group of individuals sharing major aspects of their development and their culture, and may be irrelevant to individuals of the same age developing in a different culture or very different physical environments, e.g. stone age environments vs 21st century industrialised communities. For example, tests on numerical competences in development may lead to false theories about genes for numerical competences because the results depend partly on earlier environmental influences, some of which are results of cultural evolution of the community.

This may also have serious implications for experiments on artificial genetic modification that don't take account of the complexity of epigenesis and the context-sensitivity of empirical discoveries.

The following references are to Kant's ideas about mathematical competences, and some related theories about development: Sloman(1962) Sloman on Kant(Dec 2018-work in progress), Sloman(2013c), and recent work on compositionality in biology: Sloman(2018), Sloman(2018,work in progress) on aspects of Kant's Philosophy of mathematics

This multi-stage partly abstractly-controlled form of development should also be a feature of design of future robots that have to be placed in their target environments in an under-developed state, enabling important features of their development to be influenced by inserting parameters derived from the environment into generic frameworks already in their design.

That will require robots to have far more hard-wired generic, but parametrisable, forms of intelligence than current machines have. In some cases, individual development will require several layers of creation of new abstractions formed by combining previously evolved or developed abstractions applied to new complex parameters acquired either directly from the environment, or from earlier results of this type of instantiation process. The most spectacular known example of such a process is development of language in young humans, especially processes discovered in the case of deaf children in Nicaragua, mentioned below.

Similar considerations are relevant to the design of intelligent robots.

Biological evolution involves mostly mindless processes of design, design evolution, and design instantiation combined with individual development, all in interaction with environments that can influence the developmental trajectories on different time-scales.

Much of that process involves (cooperative, competitive or neutral) interactions with both conspecifics and members of other species. If previous evolution, combined with individual learning, has built enough abstract parametrised potential solutions into individual designs, then instantiation of abstract features of evolutionary history can enormously speed up development and learning by individuals, especially if the evolutionary history has captured some immutable reusable features of spatial structures and processes.

This is related to evolution of human mathematical abilities, extending Kant's ideas Kant (1781), summarised in Sloman, Dec 2018 (work in progress), and Sloman(2018,compositionality).

Related work

The discussion above extends both ideas in Jablonka and Lamb (2005) and ideas of Chomsky and other linguists concerning the influence of prior evolution on individual language development.

I think it is also consistent with Immanuel Kant's ideas about non-empirical knowledge of synthetic necessary truths, though that connection is far from obvious. [REF NEEDED] 1962

There are links with ideas about development and learning in Karmiloff-Smith(1992), including her ideas about "representational redescription" occurring as a type of competence develops. However there are also some deep differences, as I'll try to explain later.

Relevant recent evidence regarding epigenetic mechanisms is discussed by neuroscientist Seth Grant, interviewed here by Ginger Campbell:
http://brainsciencepodcast.com/bsp/137-seth-grant
(I am grateful to Luc Beaudoin for the link.)

Bresslof's work
(added 18 Nov 2017)

I have just stumbled across the work of Paul C Bressloff Department of Mathematics, University of Utah, in particular this recent review: "Stochastic switching in biology: from genotype to phenotype" Bressloff(2017). I have not yet taken in all the details but it is highly relevant to the theory of a meta-configured genome, presented below.

Bressloff focuses on mathematical features of gene expression described at a level of detail that does not bring out some of the kinds of structure in epigenetic processes mentioned below, that seem to be much harder to investigate empirically, for example multi-layered highly structured forms of gene expression that involve instantiation of structural patterns with particular (often structured) parameters, as in formation of grammars, semantic categories, theories about space (geometry and topology), individually developed ontologies, values, preferences, and standards, and much more. Here are some samples from the paper that seem to be special cases of the general points I wish to make:

> Page 4: "The presence of transcription factors means that cellular processes can be controlled by extremely complex gene networks with multiple negative and positive feedback loops. Identifying functional modules and motifs within such networks is a major component of systems biology [2]. Gene regulation plays an essential role in viruses and bacteria, since it increases the versatility and adaptability of an organism within an environment. In multicellular organisms, gene regulation drives cellular differentiation and morphogenesis in the embryo, producing different cell types with different gene expression profiles from the same genome sequence."

> Page 82: 5.3. Metastability in epigenetics
> Aurell et al [21, 22] have developed a theoretical framework for studying the metastability of epigenetic states (see also [230]). Epigenetics concerns phenotypic states that are not encoded as genes, but as inherited patterns of gene expression originating from environmental

factors. Examples of environmental influences range from changes in the supply of nutrients in bacteria to stress in humans.

Page 127 (In section 9)
One of the major challenges in furthering our understanding of biological switches, as well as other processes in systems biology, is dealing with the complexity of most gene regulatory and biochemical signalling networks. For example, typical reaction networks involve multiple nodes (complexes) and links (reactions) that represent a hierarchy of nonlinear feedback loops. Identifying characteristic modules and motifs within these networks (including switches) is a subject of intense research. The theory of chemical reaction networks touched upon in section 2.3 could be important in identifying non-equilibrium steady-states in complex networks, as could multi-scale analyses that exploit any separation of time-scales. Another source of complexity arises when spatial effects become important.

My main concern about this is whether human brains have the power to understand some of the kinds of complexity involved in the most complex forms of gene expression producing adult human brains, e.g. the interactions between genes in an embryo and the multiple levels of gene expression that in some ancient mathematicians produced major discoveries in geometry, topology and number theory. This paper is part of an attempt to make the process more tractable by offering a theory that abstracts from many of the more complex details.

The meta-configured, multi-layered genome

The rest of this document presents ideas developed with Jackie Chappell adding more substance to the general points made so far. Biological evolution involves interacting processes at many levels of spatial scale and time scale. This document focuses mainly on a subset of processes within individual development (epigenetic processes), though many other factors have direct or indirect influences during species evolution and individual development.

The more complex the members of a species are, the more opportunities there are for individuals to benefit from processes of instantiation of parametrised generic features, thereby avoiding lengthy empirical learning. But the more complex they are the more opportunities evolution will have had to acquire such re-usable, parametrisable, design features, that enable individuals to bypass lengthy empirical learning.

A possible name for the Chappell-Sloman theory is "The Meta-Configured Genome" theory (producing a meta-configured epigenetic landscape?). The main ideas are summarised in Figure EPI, below,

One feature of the theory is that it claims that long before there were any human mathematicians, biological evolution (the blind mathematician?) implicitly made many mathematical discoveries, which were, in effect, encoded in parametrised designs for developmental processes, physical structures and useful behaviours in genomic structures that could be used (a) across different species (b) at different stages of development within a species.

Evolutionary precursors of human mathematical abilities are a special case. But less spectacular versions seem to occur in many intelligent species, including crows, mentioned below.

Could a similar diagram represent physical/chemical evolution of the universe? Genes at the top of the diagram would have to be replaced by hitherto unknown aspects of fundamental physics playing the combined role of genome and environment. (Referred to as the Fundamental Construction Kit (FCK) in Sloman[2017].)

Epigenesis in organisms and in species

Epigenetics is the study of processes and mechanisms by which genes influence the development of individual members of a species. Our ideas differ considerably from the scheme depicted as follows in 1957 by the noted biologist, C.H. Waddington, labelled "The Epigenetic Landscape".

The main difference can be summed up as follows: on the "Meta-configured" landscape, as the ball rolls down the landscape it changes the structure of the rest of the landscape, modifying the options available thereafter (for that individual), partly on the basis of effects of its previous trajectory. Think of how the early years of learning a particular language alter language-development competences thereafter. Similarly the toys and games that a child, or other intelligent animal, learns to interact with or take part in may significantly alter the kinds of things that can be learnt, or the ease with which they can be learnt later.

But that analogy is too ill defined so we try to add more content in the discussion of Figure EPI, below, explained in some detail in the rest of the paper, though most of the details will have to be provided by future research.

Figure WAD: **Waddington's view of epigenesis: The Epigenetic Landscape**
A ball rolling (passively) down a fixed landscape
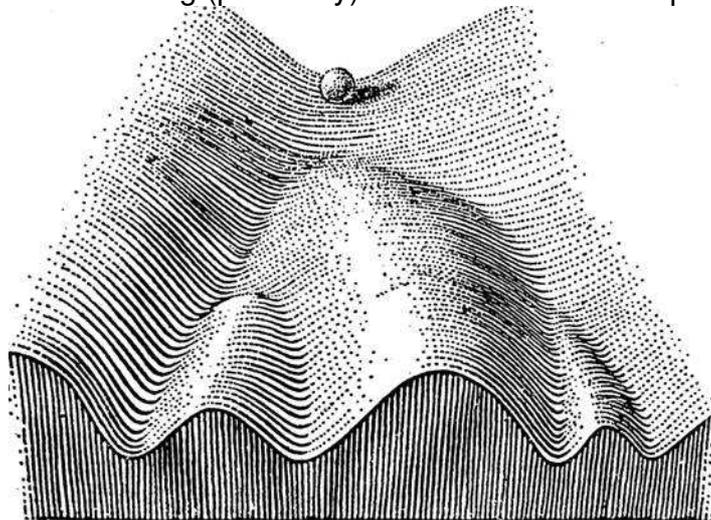
A theory of epigenesis should take account of the epigenetic requirements for a wide range of organisms, including a wide range of animals able to acquire and use information about their environments in actions varying enormously in complexity, from very simple (micro-)organisms to current highly intelligent animals such as nest building birds.

The rightmost picture in Figure Evolution below, refers (crudely) to Betty, the New Caledonian crow, who made headlines in 2002 when she unexpectedly used a straight piece of wire to make a hook in order to fish a bucket of food out of a vertical tube. What the news headlines did not say, but can be checked in videos on the Oxford Ecology lab web site, was that she made hooks in

several very different ways, all without any hesitation or trial and error, and apparently in full awareness of what she was doing and why.

Compare the discussion of toddler intelligence and weaver bird intelligence in ths 42 minute video prepared for the IJCAI workshop presentation on 19th August 2017, available here: http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/ai-cogsci-bio-sloman.webm

---

Figure Evolution
Increasing physical complexity requires increasing complexity of control
**(including increasing *dis*embodiment)**
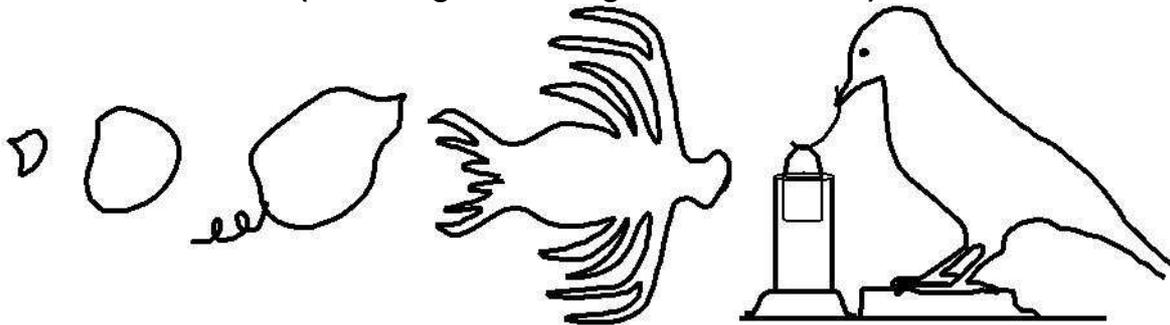
A schematic depiction of some of the transitions in biological evolution, including changes in physical form, size, habitat, capabilities and, by implication, information processing abilities required.

Many discontinuities in physical forms, behavioural capabilities, environments, types of information acquired, types of use of information and mechanisms for information-processing are still waiting to be discovered.

Introduction to the Meta-Configured Genome theory

The following figure, which extends the diagram included in Chappell and Sloman (2007) provides a crude graphical summary of some of the key ideas of the Meta-Configured Genome theory mentioned briefly above. The following sections, explaining the routes through the diagram, summarise key features of the theory. But there is still a great deal of work to be done, and the ideas presented here are likely to require both correction and extension.

Figure EPI:The meta-configured genome
A more complex theory of epigenesis (beyond Waddington)

Multiple routes from genome to behaviours
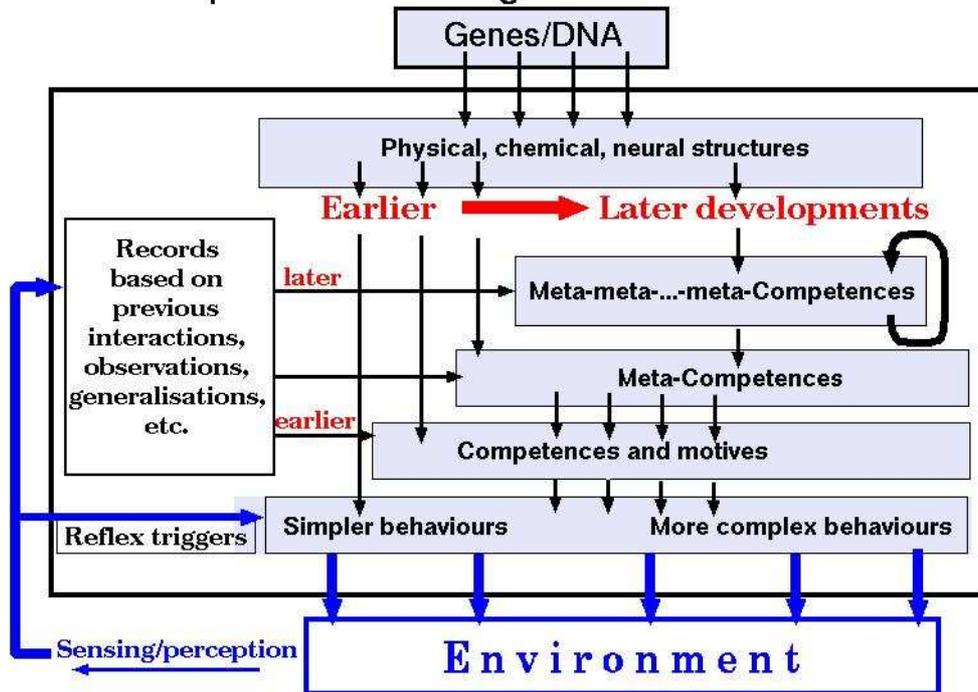
Cascaded, staggered, developmental trajectories,
with later processes using "parameters" provided by
results of earlier processes in increasingly complex ways.
Originally proposed by Chappell and Sloman (2007)
The original diagram was simpler than this one.

Added 29 Apr 2019: There is now a short (9 minute) online video presentation of the ideas in this diagram (extracted from a longer online lecture), available on youtube https://youtu.be/G8jNdBCAxVQ, and here: http://www.cs.bham.ac.uk/research/projects/cogaff/movies/#meta-config-genome

In Figure EPI, early genome-driven learning from the environment occurs in loops on the left. Downward arrows further right represent later gene-triggered processes during individual development instantiating generic patterns on the basis of results of earlier learning via feedback on left e.g. syntactic structures found, in the case of language development. Chris Miall helped with the original diagram in the Journal paper. Alan Bundy suggested the feedback loop extending meta-...competences: indicating that individual learning can expand the scope of genetic influences on acquired structures during epigenesis.

This diagram (extending the diagram in our 2007 paper) crudely summarises processes associated with a meta-configured genome. Downward arrows indicate multiple processes of gene-expression, starting with early relatively direct gene expression (down arrows on the left) and later gene expression (down arrows on the right), where later genes have parameters/gaps filled using information derived (in increasingly complex ways) from records produced by interactions with the environment during earlier gene expression (represented by down arrows more to the left).

15

The diagram does not indicate the important fact that some of the records of gene expression will be in the environment rather than in the brain. Because of that, gene expression processes in different individuals can be mutually enhancing. (There are enormously complex details involved in the processes of gene expression that must allow far more variation than suggested here.)

Restricting interactions with the environment to depend on expected (positive or negative) rewards would tremendously constrain what a learner could attempt to do, so we reject the assumption that motivation is necessarily reward-based, and instead postulate "architecture-based-motivation" (perhaps "architecture-driven" would be more appropriate). See Sloman(2009).
    See
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html,
    extending an earlier version published in 2009 in the *Newsletter on Philosophy and Computers*, (American Philosophical Association).

Roughly: motive-generation mechanisms (some triggered as complex internal reflexes) that have served ancestors well are passed on and used instinctively by descendants, who have no idea why they have some of their motives and certainly could not evaluate and compare them. Motive-generators of different types will be activated at different stages of development. For example, motives connected with mating are activated relatively late, and usually the individuals have no idea why they have those motives. In general there is no requirement for genetically specified mechanisms or behaviours to be produced at birth or soon after: the relevant genes may be expressed at a much later stage.

Educational systems that ignore these delayed effects are likely to be highly sub-optimal. E.g. a teaching strategy that focuses on producing high scores on a particular set of tests could seriously interfere with later developments that build on unnoticed side effects of earlier teaching and learning processes. For instance, emphasising phonics-based early reading may produce higher scores in reading-aloud tests, while seriously interfering with development of deeper forms of text comprehension and creative thinking whose effects cannot be measured until much later.

Background: different patterns of development in organisms

Many years ago I began to wonder at two facts based on informal observation rather than scientific study:

(a) Two (or more) types of species:
Newborn infants, or new hatchlings, of some species, e.g. deer, chickens, turtles, and many invertebrates, are highly competent, very soon after birth or hatching, including being able to move around and feed themselves or even run with the herd (e.g. newborn wildebeest foals), whereas others, e.g. crows, cats, other carnivores, and apes, including humans, are hatched or born incompetent and relatively immobile, depending crucially on parent feeding and care for some time. A biologist who attended my talk at a workshop in 1998 informed me that the former were called "precocial" and the latter "altricial" species. I later learnt from Jackie Chappell, a colleague in Biosciences, that these labels were more appropriate for specific competences than for species, as explained below.

The second informal observation was that, paradoxically(?), adult members of species born less competent ("altricial" species) seem to have more sophisticated cognitive abilities than adult precocial animals that start life much smarter. E.g. adult squirrels, crows, apes and hunting mammals seem to have greater mastery over intricate spatial structures, processes and the problems they pose than horses and deer that find their way to the mother's nipple and, in some cases, run with the herd soon after birth, or chickens and ducks that extricate themselves from eggshells and feed themselves soon after, as do many insect grubs whose parents never meet their offspring.

Informally observed differences between the two types suggested that there might be some deep clues about natural intelligence, and the tradeoffs between innate information and acquired information. For example, is there great significance in the fact that for several years human brains develop in size and structure in parallel with increasingly complex physical and cognitive capabilities often involving performance of increasingly complex and varied physical actions?

Moreover, independently of those general ideas, there are very specific well known facts, such as the fact that in humans, some aspects of genetically determined morphology, and motivational and perceptual mechanisms related to reproduction, interest in and reactions to members of the opposite sex, seem to be genetically programmed to develop relatively late in life, i.e. at puberty. So expression of genes related to both physical characteristics and information-processing abilities can be delayed for several years.

There may be less obvious but equally important genetically controlled brain functions that are delayed till later, concerned with requirements for child care. Is there also delayed gene expression related to post-parental biological functions, such as grandparenting, and perhaps varieties of motivation and behaviour relevant to competitions for pack-leadership that need to be based on a lot of prior experience? Less well controlled motivational development could lead to premature production of ambitions -- e.g. before individuals have had enough time to gain relevant experience.

In short, it seems that evolution somehow "discovered" that a type of genome for a long-lived intelligent species with many kinds of individual and social behaviours can increase evolutionary success if the processes of physical development then ageing are accompanied by genetically delayed changes in cognitive mechanisms concerned with motivation, capabilities of various kinds, and relationships with conspecifics, that differ substantially across the lifespan.

These "delayed" developments could be based partly on processes of accumulation of information within the genome, over many generations, in contrast with information acquired through individual learning. But that would restrict the knowledge and competences individuals could acquire. In contrast, if the specifications for later development are parametrised, allowing their manifestations to be tailored to specific circumstances, then parameters that vary widely across habitats and cultures could be derived from aspects of the environment during individual development, and later combined with the general design specifications, at appropriate times.

For example, genetic expression of a home-building instinct could be delayed until individuals have had time both to find out about available home-building resources in the environment and to acquire the required skills. Both of those could be facilitated by motivations connected with play and imitation of adult behaviours. Is this related to the complex nest-building achievements of some

birds, such as weaver birds and crows? Perhaps also hunting and cadaver eating competences in some mammals?

Mobility vs technical advances

Initially this type of adaptability to different environments would help a species that spreads across diverse terrain with very varied materials and circumstances available for building shelters, making tools, finding food, etc.

However, if individuals also somehow acquire abilities to invent and create new home-building materials and new techniques, then that previously evolved delayed generation e.g. of home-building motives would also be applicable to use of new materials and techniques created in that location by previous generations, not just naturally occurring diversity in resources across different regions.

I suspect that many aspects of human technology, art and science depend on such mechanisms using inheritance of abstractions that can be combined with different "fillers", some of which may also be abstractions, and which may first have evolved to meet requirements of different terrains.

(Could this be a special case of the Baldwin effect, producing new inherited abstractions rather than new inherited physical features or behaviours? https://en.wikipedia.org/wiki/Baldwin_effect)

Such pre-programmed staggered development requires delayed expression of genetic features, e.g. taste preferences, or types of curiosity, or social motives, until physical changes or types of learning have occurred that are pre-requisites for the new feature to be useful.

Examples may include changes in food preferences with age, and delayed development of mating related interests, motives, and preferences.

A deeper reason for staggering development could be related to the need to match certain developmental processes to information that has been acquired about aspects of the environment, e.g. physical aspects that need to be learnt before major forms of motivation related to social functions, jobs, etc. can develop. This could explain how the same genome works for humans in ancient technologically primitive contexts, in more recent industrialised contexts, and in very recent contexts where very young children have internet connected games or toys requiring both new kinds of motivation and new forms of physical and intellectual knowledge and skill.

These cases require a genome that does not specify motives in detail, but allows a generic, genetic specification to be instantiated using locally acquired information: a crude case might be finding out what behaviours are currently admired by conspecifics, and then generating new motives to perform those behaviours. Another type of specification could used mechanisms that detect opportunities to interact with something new or to perform a new kind of action, and then trigger production of a motive to use the opportunity: an internal reflex. (The distinction between such *architecture based* and *reward based* motivation (ABM vs RBM) is explained in Sloman(2009) -- RBM presumes some generic form of reward involved in all motivation, which I claim is a myth.)

Scientific, engineering, and artistic motivation

An important type of human motive that would normally need to be delayed beyond infancy and childhood is a desire to investigate limitations in current theories about the world and to create new hypotheses worth investigation. However, specific revolutionary intellectual curiosity and rebellious motivation should not emerge before the individual has had time to acquire broad and deep knowledge about theories already developed and their strengths and weaknesses. E.g. pioneering scientific, engineering or artistic innovators normally need to start by learning a great deal about previous achievements. Similar unique novel motivations related to modifying construction practices in local technology and artistic techniques could be generated by interactions between delayed gene expression and information, or physical features, acquired by individuals.

These ideas lead to a collection of hypotheses regarding genetic mechanisms concerned with motive generation. New types of motive may be triggered through delayed genomic expression at various stages of development. How new motives produce behaviour may be shaped by motivational mechanisms that use prior experience and knowledge of the individual, that provide parameters used to produce locally relevant special cases of generic abstract genetic specifications. So the genomic content must include "gaps" to be filled by such parameters.

In each case the expressed generic ("gappy") genes would combine with what the individual has already learnt (which can vary enormously with environment in combination with earlier gene-expression trajectories in the individual).

Use of structured parameters by generic motivators

Many computational models make use only of numerical parameters, but in biological structures and behaviours the parameters may themselves be complex structures with relationships to other structures. That idea is familiar in theoretical linguistics (especially the work of Chomsky) but is fairly new to engineering: e.g. for a long time control-engineering was almost entirely based on numerical mathematics.

The engineering applications now take their most sophisticated forms in connection with software systems that use inheritance hierarchies and parametrised structures. In such contexts, parameters may be numbers, abstract symbols, procedures, complex data structures (e.g. a set of rules for a game) or even theories.

These ideas are now familiar in theoretical computer science, AI (especially symbolic AI), theoretical linguistics, and advanced software engineering, but seem to be unknown to many others who build computational models, including models of brain function, using only forms of abstraction based on numerical parameters.

The very misleading label "Object Oriented Programming" (OOP) has come to be used for the abstraction-based software design strategy, though the label unfortunately has several very different interpretations, some more sophisticated than others. What I am suggesting is that sophisticated versions of OOP were "discovered" and put to use long ago in biological evolution, though some of the not yet understood biological examples may turn out to be far more complex than anything developed by human engineers, so far.

A fairly elementary, somewhat idiosyncratic, tutorial introduction to OOP can be found here:
http://www.cs.bham.ac.uk/research/projects/poplog/teach/oop.html

The meta-configured genome theory suggests that biological evolution "discovered" (or "stumbled across"?) the value of such parametrised design specifications (as in OOP) long before human scientists and engineers did, and as a result parametrised specifications are used in genomes for many complex species, especially portions of genomes required for information processing mechanisms and control functions. Some of them will use virtual machinery implemented in physical machinery, as explained here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-func.html
All of these ideas need to be expanded in the context of the theory of evolved construction kits, under development here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

I suspect that the most important consequences of these ideas for patterns of development controlled by genomes (epigenetic processes) have not yet been understood.

Corollaries of these evolved design strategies

I suspect that these evolutionary/developmental strategies have deep implications that have not yet been understood. I'll offer two related but different conjectures here, though there is still much empirical and theoretical work to be done. (This section needs major reorganisation, and more illustrative empirical evidence.)

Hypothesis 1 (Needs to be made clearer, more precise.)

There are some brain mechanisms that should not begin development until their development can be accompanied with actions that require a certain level of size, strength and self control. The burden of natural selection can be reduced if it does not need to encode full information about required brain features into the genome because control parameters can be picked up during earlier interactions with the environment while the relevant portions of the brain are still growing.

I am not talking about acquiring statistical regularities, but fitting generic patterns to parameters, such as muscular strength, lengths of bones, angles of rotation, moments of inertia, etc. In some cases this requires relevant physical growth to have been achieved, e.g. bones, muscles, and nervous connections, along with under-developed but usable motor control mechanisms. If relevant control parameters can be acquired from physiological sensors or from the environment, then control patterns can be instantiated using those parameters, to produce effective actions, modulated appropriately as the organism changes size, weight, strength, shape, etc.
(Add ref to D'Arcy Thompson's work, Brian Goodwin, etc.)

N.B. This proposal does not require acquired parameters to be very precise, if homeostatic (negative feedback) control mechanisms manage the details of actions. Greater precision is required, however, for ballistic actions, e.g. leaping across gaps, throwing objects at targets, etc. (Ref Alain Berthoz and Philippe Rochat?)

These ideas contradict the hypothesis that in humans there is a single very powerful, uniform, learning capability that is combined with a large initially empty, uniform, information store, as suggested (seriously??) by Turing in his *Mind* 1950 paper in *Section 7 Learning Machines*, where he wrote:

(p.456) *"Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child-brain is something like a note-book as one buys it from the stationers. Rather little mechanism, and lots of blank sheets. (Mechanism and writing are from our point of view almost synonymous.) Our hope is that there is so little mechanism in the child-brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child."*

I find it hard to believe that around 1950 Turing really believed that human brains start with "so little mechanism". He was far too intelligent and too widely read.

In contrast, the facts summarised above, but not described in detail, suggest that more intelligent species may have various sorts of learning abilities that get switched on at different stages of development. Learning abilities related to use of relatively large limbs controlled by relatively strong muscles might use specialised brain mechanisms that develop only when they are needed at a relatively late stage, and could be different from mechanisms that acquire/learn strategies for control of gaze direction, which seems to develop very early in many intelligent species.

Some mammals are born blind and begin to see some time after their birth. In those cases the required visual learning mechanisms should not begin to function until other aspects of the organism have developed, e.g. retinal functions, neural connections, and muscles required for control of gaze direction and visual focusing?

There are alternative strategies used in some genomes that produce what is known as "Biological metamorphosis", e.g. insect forms that change from caterpillar to moth or butterfly via a *pupal* stage in which molecules undergo drastic genetically controlled reorganisation, as described by Ferris Jabr in *Scientific American*.

How Did Insect Metamorphosis Evolve?
https://www.scientificamerican.com/article/insect-metamorphosis-evolution/
Added 26 Dec 2019
See also these compelling clips from a BBC TV documentary on meta-morphosis by David Malone (which I have not seen unfortunately):
https://www.bbc.co.uk/programmes/p00zv0wk/clips
(I wish producers of serious science documentaries would not add insulting spurious sound effects that make them such a struggle to listen to for growing numbers of people with age-related hearing loss (presbycusis). See also:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bbc-learning.html)

There is an additional, more subtle requirement for development of some brain mechanisms to be delayed in organisms whose evolution equips them for life in widely varying environments. There can be extreme variations in the environment that are out of the control of the genome, some of which are results of differences in terrain, climate, predators, food sources, or competitors in different parts of the planet, and some of which are results of cultural differences that are produced by earlier products of evolution. In those cases, the full details of related competences should not be encoded in the genome: each individual needs to find out what is required in the environment of its birth, or hatching, or development.

It is often assumed that this requires collection of large amounts of statistical data and computing derived regularities. The alternative assumed here is that abstract possibilities for structural variation previously "discovered" by evolutionary mechanisms are implicitly encoded in the genome but with parameters of various kinds missing, and the learning process discovers those parameters. The learning process need not be totally general (as suggested in the quotation from Turing) and the parameters acquired may be structural rather than numerical in many cases (e.g. linguistic structures that work in a community).

In some cases the processes of development and learning do not directly produce structures and behaviours but instead produce construction kits (including physical construction kits and abstract construction kits) tailored to the current environment, which then create the final products, as proposed in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

Striking examples of the use of evolved construction kits and generic patterns that are instantiated differently in different environments are differences in processes of language acquisition used by humans in different parts of this planet. Differences between spoken and sign languages are among the most spectacular examples, and will be mentioned later in connection with the question: are languages learnt or created? See the section on evolution of language below.

Hypothesis 2 (Needs to be made clearer, more precise.)

Some of the later forms of learning derive from previously evolved genetic mechanisms with a very high level of abstraction that are sensitive to abstract and complex features of the development environment. So instead of later forms of learning merely having parameters that can be replaced by different constants in different environments, they may have generalised learning mechanisms or construction kits that are instantiated in more complex ways in different environments, e.g. depending on more complex and structured parameters. Compare sophisticated uses of parametric polymorphism in modern programming languages.

Moreover, the types of learning required at different stages of development are not uniform. For example, if the claims made by Chomsky and others that language development requires particular forms of information processing related to production and perception of grammatical structures, and use of language based on those structures depends on other previously developed competences, then that might be a reason for delaying growth of relevant portions of the brain until the more basic physical and information processing mechanisms had been grown and assembled into well functioning components, and some basic features of the form of language used in the infant's culture have been absorbed -- e.g. whether it's a sign language or vocal language, etc.

That argument could apply to other capabilities also. For example, if there are forms of learning that use brain mechanisms relevant to coordination of motor control in upright sitting, standing, or walking postures, then the brain mechanisms required for planning and control functions of those coordinated late developing mechanisms need not grow until they are needed. Moreover it is possible that if grown too soon they would be stimulated in inappropriate ways and learn the wrong things, which may be an irreversible process.

Provisional summary: against innate modules, active from birth

All of the above might explain, at least in part, why some altricial species have much delayed brain development, especially as regards brain mechanisms required for more sophisticated perceptual and action control functions. (Related, but different points were made in Karmiloff-Smith(1992), discussed below.)

In all these cases different forms of learning and development may be required that cannot be pre-encoded uniformly in the genome. For example, if very different sound patterns or signing patterns have developed in different cultures, then the mechanisms for utilising those patterns to express different forms of communication syntax may not be able to use a general purpose learning engine effectively. Instead, the genome may allow each individual to grow a tailored "syntax learning system" suited to the local language culture.

If other forms of development are delayed until this mechanism has been deployed, then after they have started building on linguistic accomplishments so far, attempts to learn a new language later may not produce a second-language competence that is so well integrated into the whole system.

This may also require different forms of learning to deal with differences in semantic structures and syntactic-semantic relationships that depend on both environmental influences and previously evolved cultural structures. (Musical notation is a clear example where the latter influence is dominant.)

There are also differences that relate to culture-specific games and rituals, though it is likely that the ability to cope with those was not a selecting factor in language evolution, but a consequence of the generality of evolved mechanisms.

BACK TO CONTENTS.

Some implications

In short, what needs to be learnt in different domains at different stages may require preliminary collection of materials whose analysis influences the kind of learning required to build higher level structures and functions -- achieved by creating appropriate instantiations of very general/abstract inherited competences. The abstract genetic capabilities are combined with acquired parameters (of varying complexity).

This idea, generalising claims previously made regarding language development, was developed jointly with Jackie Chappell and presented in Sloman and Chappell (2005a), after which we were invited to write a paper for a journal in which the ideas were developed further Chappell and Sloman (2007). The ideas have developed since then, adding detail summarised in Figure EPI below.

That figure summarises the key ideas in the 2007 paper and some ideas developed later. Additional work is being done on the special case of evolution and development of mathematical abilities, especially the abilities that led to (non-numerical) geometric and topological discoveries of ancient mathematicians. Some recent related work in progress can be found in:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html

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

The latter is an incomplete draft specification of a kind of (virtual) multi-membrane machine, very different from a Turing machine, that seems to be required by an organism or machine able to make the ancient discoveries made by Archimedes, Euclid, Zeno and their (mostly unknown) predecessors. I don't yet know whether such a machine could be implemented as a virtual machine on a digital computer. Neither does anyone know how it could be implemented in brain mechanisms. I think Trehub (1991) had some ideas that could be adapted for these purposes, though he was not thinking of these problems.

This multi-layered pattern of individual development driven by a genome specifying required mechanisms at high levels of abstraction, allows evolution to produce and re-use very general abstractions that can be instantiated to produce very different types of information processing, based partly on the environment of development of each individual, which in turn can be influenced partly by what earlier generations have achieved (which is most obviously the case in humans, but may also apply to some other species).

A consequence of this type of genome is that intermediate learning mechanisms are constructed by instantiating very abstract "parametrised" learning mechanisms in different ways, using different parameters picked up at earlier stages of development, that can vary enormously as environments vary. In particular, this can use "construction kits" that the individual has created, on the basis of genetically specified construction kits of greater generality produced at earlier stages of evolution. See Sloman[2017].

We call a genome of that sort a *Metaconfigured Genome* in contrast with a *Preconfigured genome* of the kind implied by Waddington's epigenetic landscape depicted below.

These mechanisms allow individual development to be deeply influenced by aspects of the environment that have been changed by their ancestors and ancestors of other species with whom they interact. Obvious examples include types of agriculture, types of urban development, types of machinery that have been found useful, forms of energy that have been used, and many more. In Sloman & Chrisley (2005) the label "Extended genotype" was proposed for this, echoing Dawkins' "Extended Phenotype".

Variations in epigenetic trajectories

The description given so far is very abstract and allows significantly different instantiations in different species, addressing different sorts of functionality and different types of design, e.g. of physical forms, behaviours, control mechanisms, reproductive mechanisms, etc.

At one extreme the reproductive process produces individuals whose genome exercises a fixed pattern of control during development, leading to 'adults' with only minor variations.

At another extreme, instead of the process of development from one stage to another being fixed in the genome, it could be created during development through the use of more than one level of design in the genome.

E.g. if there are two or more levels then results of environmental interaction at earlier levels could determine how generic potential is later instantiated at higher levels -- e.g. what sorts of verbal morphology, grammatical structure or semantic structure develop at later stages. If there are

multiple levels then what happens at each new level may be influenced partly by results of earlier developments.

If the abstract structures are not mere numeric formulae whose instantiation involves insertion of numeric values, but also allow previously acquired structures (e.g. spatial structures, process structures, grammatical structures) to be inserted at later stages of development, then the results can differ far more than if numerical values are used. The same is true if different chemical structures are inserted into the same chemical framework. This may have been one of the insights driving Schrödinger(1944).

In a species with such multi-stage development, at intermediate stages not only are there different developmental trajectories due to different environmental influences, there are also selections among the intermediate level patterns to be instantiated, so that in one environment development may include much learning concerned with protection from freezing, whereas in other environments individual species may vary more in the ways they seek water during dry seasons.

As the development of linguistic competence shows, these processes can produce new construction kits within individuals, including construction kits for creating new (to the learner) forms of representation and reasoning. The theory of "architecture-based motivation" Sloman(2009) implies that similar structural variation can occur in patterns of motivation.

Differences in adults with a shared genome can then result partly from the influence of the environment, including the social environment, in triggering instantiation of motivational patterns. E.g. one group may learn and pass on information about where the main water holes are, and in another group individuals may learn and pass on information about which plants are good sources of water. Further details within individuals may be results of different individual experiences, or complex interactions with products of previously instantiated genetic patterns.

If these conjectures are correct, patterns of individual development in a species will normally be varied because of patterns and meta-patterns picked up by earlier generations and instantiated in cascades during individual development.

So different cultures produced jointly by a genome and previous environments can produce very different expressions of the same genome, even though individuals share similar physical forms.

The main differences are in the kinds of information acquired and used, the information processing mechanisms developed, and the motivations triggered in various contexts. But some aspects of these mechanisms will be common across the species. For example, not all cultures use advanced mathematics in designing buildings, but all use previously evolved understanding of space, time and motion.

All of this implies that evolution has found how to provide rich developmental variation by allowing information gathered by young individuals not merely to select and use pre-stored design patterns, but to create new patterns by assembling fragments of information during earlier development, then using more abstract processes to construct new abstract patterns, partly shaped by the current environment, but with the power to be used in new environments.

Nicaraguan Deaf Children
Developments in culture (including language, science, engineering, mathematics, music, literature, etc.) all show such combinations of data collection and enormous creativity, including creative ontology extension e.g. the Nicaraguan deaf children reported in this video
https://www.youtube.com/watch?v=pjtioIFuNf8
See also Ann Senghas, 2005, Language Emergence: Clues from a New Bedouin Sign Language, in *Current Biology,* 15, 12, pp. R463--R465, Elsevier,
http://dx.doi.org/10.1016/j.cub.2005.06.018

Karmiloff-Smith on Representational Redescription
This is related to the type of process Annette Karmiloff-Smith called 'Representational Re-description' (RR). But our explanations of the phenomena seem to be different. The account I am offering is that genome-encoded previously acquired abstractions 'wait' to be instantiated at different stages of development, using cascading alternations between data-collection and abstraction formation (RR) by instantiating higher level generative abstractions (e.g. meta-grammars), *not by forming statistical generalisations*. Those abstractions can be thought of as evolved modules that include parameters. The parameters come from processes of individual development and influence the ways in which the evolved abstractions are instantiated.

This could account for both the great diversity of human languages and cultures, and the power of each one, all supported by a common genome operating in very different environments.

Jackie Chappell noticed the implication that instead of the genome specifying a fixed 'epigenetic landscape' (as proposed at one stage by Waddington) it provides a schematic landscape and mechanisms that allow each individual (or in same cases groups of individuals) to modify the landscape while moving down it (e.g. adding new hills, valleys, channels, barriers, taps that can be turned on or off, etc.).

A different view was proposed by Annette Karmiloff-Smith, e.g. in Karmiloff-Smith(1992) Karmiloff-Smith(2006). As I understand her, she thinks the theory of evolved (innate) modules is completely wrong. Instead she proposes that the emergence of new competences during development comes from the complex dynamics of development of neuronal structures specified in the genome, interacting in increasingly complex ways during development, partly under the influence of the environment. If such interactions can produce new competences there is no need for those competences to be specified explicitly in the genome as suggested in Fodor(1975) or to be produced by explicit training combined with a general purpose learning mechanism.

According to the Meta-Configured Genome hypothesis they are both wrong: there are many late-developing competences that are neither produced using a general purpose (e.g. reward-based) learning mechanism, nor provided innately at birth by the genome, nor produced by side-effects of complex interacting neural development processes. Instead there are mechanisms provided by the genome which (a) do not become active at a very early stage of development and (b) are not "serendipitous" results (side-effects) of interacting processes of neural development, but rather are specified at a high level of abstraction in the genome with missing parameters, and activated only at a relatively late stage of development when values for the parameters have (implicitly) been collected from interactions with the environment as side-effects of earlier gene-expression processes.

Perhaps Immanuel Kant would have accepted this explanation as an elaboration of his theory that there are kinds of knowledge that are not *derived from* experience, but are in a sense *awakened by* experience Kant(1781).

Though most visible in language development, the process is not unique to language development, but occurs throughout childhood (and beyond) in connection with many aspects of development of information processing abilities, construction of new ontologies, theory formation, etc.

This differs from forms of learning or development that use uniform statistics-based methods for repeatedly finding patterns at different levels of abstraction.

Instead, Figure EPI indicates that the genome encodes increasingly abstract and powerful creative mechanisms developed at different stages of evolution, that are 'awakened' (a notion used by Kant) in individuals only when appropriate, so that they can build on what has already been learned or created in a manner that is tailored to the current environment.

For example, in young (non-deaf) humans, processes giving sound sequences a syntactic interpretation develop after the child has learnt to produce and to distinguish some of the actual speech sounds used in that location.

In social species, the later stages of Figure EPI include mechanisms for discovering non-linguistic ontologies and facts that older members of the community have acquired, and incorporating relevant subsets in combination with new individually acquired information.

Instead of merely absorbing the details of what older members have learnt, the young can absorb forms of creative learning, reasoning and representation that older members have found useful and apply them in new environments to produce new results.

In humans, this has produced spectacular effects, especially in the last few decades.

The evolved mechanisms for representing and reasoning about possibilities, impossibilities and necessities were essential for both perception and use of affordances and for making mathematical discoveries, something statistical learning cannot achieve.

BACK TO CONTENTS.

Evolution as epigenesis

Added 18 Aug 2017
The theory of evolved construction kits mentioned above (see Sloman[2017]), can be understood as claiming that something loosely similar to the above model of epigenesis applies also to the processes of biological evolution, except that the initial state is not a product of biological evolution. The physical state of a newly formed earth-like planet, like a genome, has enormous potential for supporting "epigenetic" processes that eventually lead to billions of different forms of life.

Evolution and use of derived construction kits (DCKs), both concrete and abstract (as required for many forms of information processing), uses a process of multi-layer feedback partly like the Chappell-Sloman conjectures depicted in Figure EPI.

The repeated use by evolution of increasingly complex and abstract designs, shows that the theory of evolution by design is basically correct, except that there is no single initial design that explains everything: instead new designs are repeatedly created by evolutionary mechanisms and put to use. That seems to include designs used in ancient mathematical minds, still available in the human genome, whose precise features have so far eluded AI researchers, and everyone else.

Space-time

An invariant for all species in this universe is space-time embedding, and changing spatial relationships between body parts and things in the environment.

The relationships vary between water-dwellers, cave-dwellers, tree-dwellers, flying animals, and modern city-dwellers.

Representational requirements depend on body parts and their controllable relationships to one another and other objects.

So aeons of evolution will produce neither a tabula rasa nor geographically specific spatial information, but a collection of generic mechanisms for finding out what sorts of spatial structures have been bequeathed by ancestors as well as physics and geography, and learning to make use of whatever is available McCarthy (2008): that's the main reason why embodiment is relevant to evolved cognition. However, as organisms grow more complex, with more complex behavioural alternatives and more options to choose between, requiring actions extended over space and time, cognitive processing must become increasingly disembodied, i.e. disconnected from sensory and motor subsystems. For example, if some plan unexpectedly goes badly wrong, working out how that happened may require reflecting in detail on the processes that occurred and the un-tried alternatives. The results of that process of analysis may prove life-saving at some future time. Enactivist, anti-cognitivist, anti-computational theories of cognition are rightly given short shrift in Rescorla (2015). (Most of the achievements of human mathematics, science and engineering, and even philosophy, would have been impossible if enactivist/embodied theories of cognition were close to the truth. Developing an enactivist/disembodied theory of cognition is itself a mostly disembodied process!)

Kant's ideas about geometric knowledge are relevant though he assumed that the innate apparatus was geared only to structures in Euclidean space, whereas our space is only approximately Euclidean.

Somehow the mechanisms conjectured in Figure 2 eventually (after many generations) made it possible for humans to make the amazing discoveries recorded in Euclid's Elements, still used world-wide by scientists and engineers.

If the parallel axiom is removed what remains is still a very rich collection of facts about space and time, especially topological facts about varieties of structural change, e.g. formation of networks of relationships, deformations of surfaces, and possible trajectories constrained by fixed obstacles.

If we can identify a type of construction-kit that produces young robot minds able to develop or evaluate those ideas in varied spatial environments, we may find important clues about what is missing in current AI.

Long before logical and algebraic notations were used in mathematical proofs, evolution had produced abilities to represent and reason about what Gibson called 'affordances', including types of affordance that as far as I know he did not consider, namely reasoning about possible and impossible alterations to spatial configurations and necessary consequences of possible alterations. How brains represent possibilities, impossibilities, and necessary consequences (as opposed to learnt associations) is, as far as I know, still open.

Example:

The (topological) impossibility of solid linked rings becoming unlinked is usually obvious even to children without any mathematical education. E.g.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/rings.html

This rubber-band example is harder to understand:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/rubber-bands.html

I suspect brains of many intelligent animals make use of topological reasoning mechanisms that have so far not been discovered by brain scientists or AI researchers. Weaver birds are among the most obvious candidates.

Addition of meta-cognitive mechanisms able to inspect and experiment with reasoning processes may have led both to enhanced spatial intelligence and meta-cognition, and also to meta-metacognitive reasoning about other intelligent individuals.

BACK TO CONTENTS.

Other species

I suspect that further investigation will reveal varieties of information processing (computation) that have so far escaped the attention of researchers, but which play important roles in many intelligent species, including not only humans and apes but also elephants, corvids, squirrels, cetaceans and others.

In particular, some intelligent non-human animals and pre-verbal human toddlers seem to be able to use mathematical structures and relationships (e.g. partial orderings and topological relationships) unwittingly. Mathematical meta-meta...-cognition seems to be restricted to humans, but develops in stages, as Piaget found (1952). E.g. it seems that recognition of transitivity of one-one correspondence is not achieved until about the 5th or 6th year. (What would have to change in brains to make that occur?)

However, I suspect that (as Kant seems to have realised) the genetically provided mathematical powers of intelligent animals make more use of topological and semi-metrical geometric reasoning, using analogical, non-Fregean, representations, than the logical, algebraic, and statistical capabilities that have so far dominated AI and robotics. (Sloman (1971) and chapter 7 of Sloman 1978)

For example, even the concepts of cardinal and ordinal number are crucially related to concepts of one-one correspondence between components of structures, most naturally understood as a topological relationship rather than a logically definable relationship.
http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap8.html

(NB 'analogical' does not imply 'isomorphic' as often suggested. A typical 2D picture (an analogical representation) of a 3D scene cannot be isomorphic with the scene depicted. A 3D to 2D projection is not an isomorphism except in special cases. There is a deeper distinction between Fregean and Analogical forms of representation Sloman (1971), concerned with the relationships between representation and what is represented.

BACK TO CONTENTS.

Disembodiment of cognition evolves (Epigenesis of evolutionary mechanisms}

All this shows why increasing complexity of physical structures and capabilities, providing richer collections of alternatives and more complex internal and external action-selection criteria, requires increasing disembodiment of information processing -- as the information required for particular episodes of control, motive generation, or planning becomes more abstract, more concerned with remote locations more concerned with past events or future possibilities, less concerned with reference to physical objects, actions and other processes, and more concerned with things like past percepts, memories, possible future decisions, modes of reasoning, and mental states of other individuals (what they perceive, want, intend, know, can do, etc.)

Such transitions occur both in individual development in intelligent species and also in evolution of complex organisms.

The fact that evolution is not stuck with the Fundamental Construction Kit (FCK) provided by physics and chemistry, but also produces and uses new 'derived' construction-kits (DCKs), including abstract construction kits needed for intelligent organisms (e.g. grammar construction kits in humans), enhances both the mathematical and the ontological creativity of evolution, which is indirectly responsible for all the other known types of creativity.

Although I have not developed the idea in this paper, the work on construction kits and their **essential** role in evolution on this planet, suggests that there are weak but important analogies between epigenetic processes in individual humans illustrated in Figure EPI, and some evolutionary processes. In both cases, the development depends on discovery of powerful abstractions ("moving upwards") that can be instantiated in different ways in different species or the same species at different times ("moving downwards"), instead of all evolution being simply "sideways" movement at a fixed level of abstraction in design.

(This distinction is ignored by theories of mathematical discovery by humans that emphasise use of metaphor and analogy instead of use of abstraction and multiple re-instantiation.)

The fact that many evolved construction kits, and their products, depended on natural selection "discovering" enormously powerful re-usable mathematical abstractions, whose re-use involved not just copying, but instantiation of a generic schema with new parameters, illustrates a partial analogy between epigenesis in intelligent organisms and epigenesis in evolution. To that extent I am proposing that some aspects of evolution need intelligent design, more precisely intelligent abstraction of powerful re-usable structures from special cases, but all the intelligence used in evolutionary design processes was previously produced by evolution.

This helps to counter both the view that all of mathematics is a product of human minds, and a view of metaphysics as being concerned with something unchangeable.

The notion of 'Descriptive Metaphysics' presented by Strawson (1959) needs to be revised, to include 'Meta-Descriptive Metaphysics'.

BACK TO CONTENTS.

---

and operate in ways that are informed by results of previous learning, and therefore should not be turned on at birth, even if there are mechanisms available that could do so.

This contrasts with the once-popular theory that humans are born so helpless because if they were retained in the mother until their brains were fully formed the mothers would be more likely to die in childbirth; or the alternative theory that walking upright is not compatible with having a large enough pelvic aperture for a birth several years later. A different explanation is offered in this Scientific American blog:

"Rosenberg additionally noted--and I found this especially fascinating--that the authors mention the possibility that the timing of birth actually optimises cognitive and motor neuronal development. That idea, first proposed by Swiss zoologist Adolf Portman in the 1960s, is worth pursuing, she says. 'Maybe human newborns are adapted to soaking up all this cultural stuff and maybe being born earlier lets you do this,' she muses. 'Maybe being born earlier is better if you're a cultural animal.' Food for thought."
https://blogs.scientificamerican.com/observations/why-humans-give-birth-to-helpless-babies/

Types of competence, rather than types of species
When Jackie Chappell came to the University of Birmingham in 2004 we began to discuss these ideas and she quickly pointed out that it was incorrect to assume that altricial species were uniformly incompetent at birth. For example, most are very competent at sucking soon after birth. She suggested that it would be more appropriate to contrast precocial and altricial competences, allowing that species may differ in the initial variety and proportions of competences of both sorts. She also helped me see that future intelligent robots would require similar mixtures -- including precocial competences used for bootstrapping and altricial competences that begin to develop only after information gained by using precocial competences has been acquired, analysed and organised. This also allowed some species with a long process of development to have genomes that initiated development of more sophisticated altricial competences at different stages of development.

In 2005, we presented a paper suggesting that future intelligent robots would also require such layers of development in which more sophisticated types of competence could begin their development at later stages, building on earlier competences.

To be added

The Super-Turing membrane machine, for perception and spatial reasoning.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-geom.html

REFERENCES AND LINKS
(Temporary list: to be pruned)

PAPERS WITH JACKIE CHAPPELL

Jackie Chappell and Aaron Sloman, (2007) Natural and artificial metaconfigured altricial information-processing systems, *International Journal of Unconventional Computing,* 3(3), 221-239.
http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717

Jackie Chappell and Aaron Sloman, Two ways of understanding causation: Humean and Kantian, Invited contributions to *WONAC: International Workshop on Natural and Artificial Cognition* Pembroke College, Oxford, June 25-26, 2007,
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac

Aaron Sloman, 2010, The Design-Based Approach to the Study of Mind (in humans, other animals, and machines) Including the Study of Behaviour Involving Mental Processes, *Proc. Int. Symposium on AI-Inspired Biology (AIIB), AISB 2010 convention,* Eds. J. Chappell, N. A. Hawes, S. Thorpe and A. Sloman, pp. 87--94, De Montfort University, Leicester,
http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1002

Aaron Sloman, Genomes for self-constructing, self-modifying information-processing architectures, Invited talk at *SGAI 2010 Workshop on Bio-inspired and Bio-Plausible Cognitive Robotics,* Cambridge, 2010, Dec, http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk89

OTHER REFERENCES

Graham Bell, *Selection The Mechanism of Evolution,* OUP, 2008. Second Edition.

Alexander Bogomolny, (2017) Pythagorean Theorem and its many proofs from *Interactive Mathematics Miscellany and Puzzles*
http://www.cut-the-knot.org/pythagoras/index.shtml

Paul C Bressloff, 2017. Stochastic switching in biology: from genotype to phenotype, *Journal of Physics A: Mathematical and Theoretical* 50 (2017)
https://doi.org/10.1088/1751-8121/aa5db4

Euclid and John Casey, *The First Six Books of the Elements of Euclid,* Project Gutenberg, Salt Lake City, Apr, 2007,
http://www.gutenberg.org/ebooks/21076
Also see "The geometry applet"
   http://aleph0.clarku.edu/~djoyce/java/elements/toc.html (HTML and PDF)

N. Chomsky, 1965, *Aspects of the theory of syntax,* MIT Press, Cambridge, MA.

Shang-Ching Chou, Xiao-Shan Gao and Jing-Zhong Zhang, 1994, *Machine Proofs In Geometry: Automated Production of Readable Proofs for Geometry Theorems,* World Scientific, Singapore, http://www.mmrc.iss.ac.cn/~xgao/paper/book-area.pdf

Juliet C. Coates, Laura A. Moody, and Younousse Saidi, "Plants and the Earth system - past events and future challenges', *New Phytologist,* 189, 370-373, (2011).

Alan Turing - His Work and Impact, eds., S. B. Cooper and J. van Leeuwen, Elsevier, Amsterdam, 2013. (contents list).

Kenneth Craik, 1943, *The Nature of Explanation,* Cambridge University Press, London, New York,

Deaf Studies Trust, 'Deaf Children Beginning to Sign', A paper prepared for Frank Barnes School, London, 17th February 1995

D. C. Dennett, 1978 *Brainstorms: Philosophical Essays on Mind and Psychology*. MIT Press, Cambridge, MA.

D. C. Dennett, 1995, *Darwin's Dangerous Idea: Evolution and the Meanings of Life,* Penguin Press, London and New York,

D.C. Dennett, 1996
Kinds of minds: towards an understanding of consciousness,
Weidenfeld and Nicholson, London, 1996,
http://www.amazon.com/Kinds-Minds-Understanding-Consciousness-Science/dp/0465073514

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

J.A. Fodor, 1983, *The Modularity of Mind* MIT Press, Cambridge MA.

T. Froese, N. Virgo, and T. Ikegami, Motility at the origin of life: Its characterization and a model', *Artificial Life,* 20(1), 55-76, (2014).

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

Gallistel, C.R. & Matzel, L.D., 2012(Epub), The neuroscience of learning: beyond the Hebbian synapse, *Annual Revue of Psychology,* Vol 64, pp. 169--200, https://doi.org/10.1146/annurev-psych-113011-143807

H. Gelernter, 1964, Realization of a geometry-theorem proving machine, in *Computers and Thought,* Eds. Feigenbaum, Edward A. and Feldman, Julian, pp. 134-152, McGraw-Hill, New York, Re-published 1995 (ISBN 0-262-56092-5),
http://dl.acm.org/citation.cfm?id=216408.216418
Also at
https://pdfs.semanticscholar.org/2edc/8083073837564306943aab77d6dcc19d0cdc.pdf

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

Seth G.N. Grant, 2010, Computing behaviour in complex synapses: Synapse proteome complexity and the evolution of behaviour and disease, *Biochemist* 32, pp. 6-9,
http://www.biochemist.org/bio/default.htm?VOL=32&ISSUE=2

M. M. Hanczyc and T. Ikegami, 'Chemical basis for minimal cognition', *Artificial Life,* 16, 233-243, (2010).

John Heslop-Harrison, *New concepts in flowering-plant taxonomy,* Heinemann, London, 1953.

Peter Hoffmann (video lecture), Life's Ratchet: How Molecular Machines Extract Order from Chaos, November 19, 2012,
https://www.microsoft.com/en-us/research/video/lifes-ratchet-how-molecular-machines-extract-order-from-chaos/

Peter M Hoffmann, 2016, How molecular motors extract order from chaos (a key issues review) 10 February 2016, *Reports on Progress in Physics,* Volume 79, Number 3, IOP Publishing Ltd.
https://iopscience.iop.org/article/10.1088/0034-4885/79/3/032601/meta

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

Immanuel Kant, *Critique of Pure Reason,* Macmillan, London, 1781. Translated (1929) by Norman Kemp Smith.
Various online versions are also available now.

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

A. Karmiloff-Smith, 2006, The tortuous route from genes to behavior: A neuroconstructivist approach, *Cognitive, affective & behavioral neuroscience,* vol 6, pp. 9--17,
https://doi.org/10.3758/CABN.6.1.9

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

M.W. Kirschner and J.C. Gerhart, *The Plausibility of Life: Resolving Darwin's Dilemma,* Yale University Press, Princeton, 2005.

D. Kirsh, "Today the earwig, tomorrow man?', *Artificial Intelligence,* 47(1), 161-184, (1991).

I. Lakatos, 1976, *Proofs and Refutations,* Cambridge University Press, Cambridge, UK,

John McCarthy and Patrick J. Hayes, 1969, "Some philosophical problems from the standpoint of AI", *Machine Intelligence 4,* Eds. B. Meltzer and D. Michie, pp. 463--502, Edinburgh University Press,
http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html

John McCarthy, "The well-designed child', *Artificial Intelligence,* 172(18), 2003-2014, (2008). (Written in 1996).

Tom McClelland, (2017) AI and affordances for mental action, in Computing and Philosophy Symposium, *Proceedings of the AISB Annual Convention 2017* pp. 372-379. April 2017. http://wrap.warwick.ac.uk/87246

Tom McClelland, 2017 The Mental Affordance Hypothesis" MindsOnline 2017 http://mindsonline.philosophyofbrains.com/2017/2017-session-1/the-mental-affordance-hypothesis/ Video presentation https://www.youtube.com/watch?v=zBqGC4THzqg

Nathaniel Miller, 2007, *Euclid and His Twentieth Century Rivals: Diagrams in the Logic of Euclidean Geometry,* Center for the Study of Language and Information, Stanford Studies in the Theory and Applications of Diagrams, https://web.stanford.edu/group/cslipublications/cslipublications/site/9781575865072.shtml

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

Jean Piaget, 1952, *The Child's Conception of Number,* Routledge \& Kegan Paul, London.

W. T. Powers, *Behavior, the Control of Perception,* Aldine de Gruyter, New York, 1973.

Michael Rescorla, (2015) The Computational Theory of Mind, in *The Stanford Encyclopedia of Philosophy,* Ed. E. N. Zalta, Winter 2015, http://plato.stanford.edu/archives/win2015/entries/computational-mind/

Philip Robbins (2017) "Modularity of Mind", *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/win2017/entries/modularity-mind/

Philippe Rochat, 2001, *The Infant's World,* Harvard University Press, Cambridge, MA,

C. Rutz, S. Sugasawa, J E M van der Wal, B C Klump, & J St Clair, 2016, 'Tool bending in New Caledonian crows' in *Royal Society Open Science*, Vol 3, No. 8, 160439. DOI: 10.1098/rsos.160439 https://research-repository.st-andrews.ac.uk/handle/10023/9283

A. Sakharov (2003 onwards) *Foundations of Mathematics (Online References)* Alexander Sakharov, with contributions by Bhupinder Anand, Harvey Friedman, Haim Gaifman, Vladik Kreinovich, Victor Makarov, Grigori Mints, Karlis Pdnieks, Panu Raatikainen, Stephen Simpson, "This is an online resource center for materials that relate to foundations of mathematics (FOM). It is intended to be a textbook for studying the subject and a comprehensive reference. As a result of this encyclopedic focus, materials devoted to advanced research topics are not included. The author has made his best effort to select quality materials on www." http://sakharov.net/foundation.html
NOTE: some of the links to other researchers' web pages are out of date, but in most cases a

search engine should take you to the new location.

Dana Scott, 2014, Geometry without points. (Video lecture, 23 June 2014,University of Edinburgh)
https://www.youtube.com/watch?v=sDGnE8eja5o

Erwin Schrödinger, *What is life?,* CUP, Cambridge, 1944.
Commented extracts available here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html

Claude Shannon, (1948), A mathematical theory of communication, in *Bell System Technical Journal,* July and October, vol 27, pp. 379--423 and 623--656,
https://archive.org/download/pdfy-nl-WZBa8gJFl8QNh/shannon1948.pdf

Stewart Shapiro, 2009 We hold these truths to be self-evident: But what do we mean by that? *The Review of Symbolic Logic*, Vol. 2, No. 1
https://doi.org/10.1017/S175502030909011X

A. Sloman, 1962, *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth* (DPhil Thesis), PhD. dissertation, Oxford University, (now online)
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-1962

A. Sloman, 1971, "Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence", in *Proc 2nd IJCAI,* pp. 209--226, London. William Kaufmann. Reprinted in *Artificial Intelligence,* vol 2, 3-4, pp 209-225, 1971.
http://www.cs.bham.ac.uk/research/cogaff/62-80.html#1971-02
An expanded version was published as chapter 7 of Sloman 1978, available here.

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

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

A. Sloman, 1996, Actual Possibilities, in *Principles of Knowledge Representation and Reasoning* (Proc. 5th Int. Conf on Knowledge Representation (KR '96)), Eds. L.C. Aiello and S.C. Shapiro, Morgan Kaufmann, Boston, MA, pp. 627--638,
http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15

A. Sloman, (2000) "Interacting trajectories in design space and niche space: A philosopher speculates about evolution', in *Parallel Problem Solving from Nature (PPSN VI),* eds. M.Schoenauer, et al. Lecture Notes in Computer Science, No 1917, pp. 3-16, Berlin, (2000). Springer-Verlag.

A. Sloman, 2001, Evolvable biologically plausible visual architectures, in *Proceedings of British Machine Vision Conference,* Ed. T. Cootes and C. Taylor, BMVA, Manchester, pp. 313--322,
http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#76
http://dx.doi.org/10.5244/C.15.33},

A. Sloman, 2002, The irrelevance of Turing machines to AI, in *Computationalism: New Directions,* Ed. M. Scheutz, MIT Press, Cambridge, MA, pp. 87--127,
http://www.cs.bham.ac.uk/research/cogaff/00-02.html#77

A. Sloman and R.L. Chrisley, (2003) "Virtual machines and consciousness', *Journal of Consciousness Studies,* 10(4-5), 113-172.

A. Sloman and R. L. Chrisley, 2005, More things than are dreamt of in your biology: Information-processing in biologically-inspired robots, *Cognitive Systems Research,*, 6, 2, June 2005. pp. 145--174
http://www.cs.bham.ac.uk/research/projects/cogaff/04.html#200408

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

A. Sloman and J. Chappell, 2005b, Altricial self-organising information-processing systems, *AISB Quarterly 121*, Summer 2005, pp. 5--7,
http://www.cs.bham.ac.uk/research/cogaff/05.html#200503,

A. Sloman, 2008, The Well-Designed Young Mathematician, *Artificial Intelligence,* 172, 18, pp. 2015--2034, Elsevier.
http://www.cs.bham.ac.uk/research/projects/cogaff/08.html#807

A. Sloman, (2008a). Architectural and representational requirements for seeing processes, proto-affordances and affordances. In A. G. Cohn, D. C. Hogg, R. Moller, & B. Neumann (Eds.), *Logic and probability for scene interpretation.* Dagstuhl, Germany: Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.
http://drops.dagstuhl.de/opus/volltexte/2008/1656

A. Sloman, 2009-2019, Architecture-Based Motivation vs Reward-Based Motivation, *Newsletter on Philosophy and Computers,* American Philosophical Association, 09,1, pp. 10--13, Newark, DE, USA. (This version has been extended since publication.)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html

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

A.Sloman (2010-2017), A (Possibly) new kind of (non?) Euclidean geometry, based on an idea by Mary Pardoe.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/p-geometry.html

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

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

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

Aaron Sloman (2007-2014)
Unpublished discussion Paper: Predicting Affordance Changes: Steps towards knowledge-based visual servoing. (Including videos).
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html

A. Sloman (2015). What are the functions of vision? How did human language evolve? Online research presentation.
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111

A. Sloman 2017, "Construction kits for evolving life (Including evolving minds and mathematical abilities.)" Technical report (work in progress)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

(An earlier version, frozen during 2016, was published in a Springer Collection in 2017:
https://link.springer.com/chapter/10.1007%2F978-3-319-43669-2_14
in *The Incomputable Journeys Beyond the Turing Barrier*
Eds: S. Barry Cooper and Mariya I. Soskova
https://link.springer.com/book/10.1007/978-3-319-43669-2)

Aaron Sloman, Jackie Chappell and the CoSy PlayMate team, 2006, Orthogonal recombinable competences acquired by altricial species (Blankets, string, and plywood) School of Computer Science, University of Birmingham, Research Note COSY-DP-0601,
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/orthogonal-competences.html

Aaron Sloman and David Vernon. A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots, 2007. Contribution to euCognition wiki.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-requirements.html

Aaron Sloman, 2018(comp),
Biologically Evolved Forms of Compositionality: Structural relations and constraints *vs* Statistical correlations and probabilities
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html
(Expanded version of paper accepted for Syco-1 workshop on compositionality
http://events.cs.bham.ac.uk/syco/1/ September 2018)
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.pdf

Aaron Sloman, Dec 2018 (work in progress), Key Aspects of Immanuel Kant's Philosophy of mathematics ignored by most psychologists and neuroscientists studying mathematical competences.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html

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

Max Tegmark, 2014, *Our mathematical universe, my quest for the ultimate nature of reality,* Knopf (USA) Allen Lane (UK), (ISBN 978-0307599803/978-1846144769)

Arnold Trehub, 1991, *The Cognitive Brain,* MIT Press, Cambridge, MA,
http://people.umass.edu/trehub/

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

A. M. Turing, "Computing machinery and intelligence', *Mind*, 59, 433-460, (1950). (reprinted in E.A. Feigenbaum and J. Feldman (eds) Computers and Thought McGraw-Hill, New York, 1963, 11-35).

A. M. Turing, (1952) "The Chemical Basis Of Morphogenesis", *Phil. Trans. Royal Soc. London B* 237, 237, 37-72.

Note: A presentation of Turing's main ideas for non-mathematicians can be found in
Philip Ball, 2015, "Forging patterns and making waves from biology to geology: a commentary on Turing (1952) 'The chemical basis of morphogenesis'",
http://dx.doi.org/10.1098/rstb.2014.0218

Barbara Vetter (2011), Recent Work: Modality without Possible Worlds, *Analysis,* 71, 4, pp. 742--754,
https://doi.org/10.1093/analys/anr077

C. H. Waddington, 1957 *The Strategy of the Genes. A Discussion of Some Aspects of Theoretical Biology,* George Allen & Unwin, 1957.

R. A. Watson and E. Szathmary, "How can evolution learn?', *Trends in Ecology and Evolution,* 31(2), 147-157, (2016).

Weir, A A S and Chappell, J and Kacelnik, A, (2002) Shaping of hooks in New Caledonian crows, *Science,* vol 297, p 981,

L. Wittgenstein (1956), *Remarks on the Foundations of Mathematics*, translated from German by G.E.M. Anscombe, edited by G.H. von Wright and Rush Rhees, first published in 1956, by Blackwell, Oxford. There are later editions. (1978: VII 33, p. 399)

About this document
Revised versions of this document will be available here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html
and a PDF version.

-

-
-
-

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

Creative Commons Licence