

NB: UNFINISHED DRAFT
(TEMPORARY VERSION)

Below are revised versions of figures used since June 2022 for talks entitled:
Recently hatched, and depicted, ideas about hatching and intelligence[**]

Using very low energy physics and chemistry
at "normal" temperatures
in egg-laying vertebrates

By Aaron Sloman
<http://www.cs.bham.ac.uk/~axs>
Emeritus Professor of Artificial Intelligence and Cognitive Science.
[School of Computer Science](#)
[University of Birmingham, UK](#)

Note:

There are likely to be serious gaps or errors in my reports and conjectures below, because this work relates to a very wide range of research fields. I'll be grateful for information about such flaws, and suggestions for correcting them.

This Document

This document is available at:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-devo-figs.html>

I can provide a PDF version if requested.

NB THIS WEB PAGE IS STILL UNDER CONSTRUCTION AND LIKELY TO CHANGE

Some of the changes are listed here. Last updated (04 Nov 2022).

23 Oct 2022: Quite a lot of reorganisation. Added references.

[Revised version of Evo-Devo diagram](#)

17 Oct 2022: added links to [Related work](#)

15 Sep 2022: New versions of diagrams: development going left to right and evolution top to bottom, instead of the other way round in [the earlier version](#).

10 & 17 Sep 2022: Added information about [meta-configured genomes](#).

9 Sep 2022: Added [link to avocet video](#).

[] BACKGROUND: About the title ("Recently Hatched ...")**

This could be thought of as a talk on philosophy of physics: one of the many examples of the fact there can be deep overlap between theories about X and philosophy of X. But I have not encountered philosophers of science who are interested in paying attention to the sorts of chemical/biological phenomena described below, although there are physicists who have been interested, e.g. Schroedinger in his *What is life?* ([1944](#))

In October 2020 I paid attention for the first time to the fact that well-known competences of newly hatched animals, especially vertebrate hatchlings, pose serious challenges for current theories in biology, neuroscience, psychology, philosophy and related disciplines, and potentially also AI.

During 2021, as the ideas developed, I gave several invited talks presenting facts and speculations about hatching processes in (vertebrate) eggs, and their relevance to various research areas, including philosophy of mathematics, epistemology, philosophy of mind and philosophy of biology, as well as developmental biology and neuroscience.

During 2022 the ideas about hatching processes grew increasingly complex as I noticed more details. The new ideas were presented in zoom talks starting in March 2022. The ideas continued to develop between talks. This web page was set up in October 2022 in an attempt to present the main ideas as compactly as possible, combining portions of previous web pages used for presentations. New developments will be accommodated as they occur.

Note:

This slightly shorter title "Recently hatched, ideas about hatching and intelligence", without "and depicted" was used for talks given since March 2022, though the contents of the talks changed as my understanding of the problems and possible solutions changed -- an ongoing process.

The phrase "and depicted" was added to the title above on 25 Oct 2022.

Those presentations used various subsets of a large, messy, constantly growing web page, which in turn is part of a very large, messy and frequently expanded web site with many presentations linking aspects of philosophy, mathematics, biology, chemistry, physics, computer science, AI, psychology, education, ...

[with a lot of references:](#)

- to earlier work on this project,
- to related online or published information, and
- to a variety of disputes about the nature of mathematical discovery, and the cognitive/biological mechanisms that made possible ancient mathematical discoveries, including discoveries concerning geometry and topology.

This is a new version of an older, longer, fairly indigestible web site:

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

This shorter version of the "hatching-talks" web site was created in October 2022 and, with luck, will survive for some time, with corrections and extensions. (It is no longer as short as it once was!)

Background information that took up a lot of space near the top of earlier versions has either been moved down the page or replaced by links to separate online documents.

For a while, during 2022, my talks used a single monolithic Evo-Devo diagram [below](#), inspired by a pencil sketch <https://www.cs.bham.ac.uk/~axs/fig/stepney-fig.jpg> received from Professor Susan Stepney (<https://www.cs.york.ac.uk/people/susan>) after she heard me give a talk about hatching processes early in 2022.

Inspired, at first, by her sketch, I started using more complex diagrams presented [below](#) emphasising different aspects of a large, still incomplete, collection of ideas linking mechanisms and processes of evolution and development, and their consequences.

The remainder of this document omits much of the background detail that was included in earlier versions that eventually grew too long and too disorganised, e.g.:

<https://www.cs.bham.ac.uk/research/projects/cogaff/misc/hatching-talks-2022.html>

This new document/presentation includes only a key subset of the ideas. I hope the ideas are easier to understand/digest than a full description of would be. I am not claiming that research in biology has so far provided details replies to questions about addition to being much shorter (for now!).

The contrast between development of a foetus inside an eggshell and development in a uterus (e.g. in mammals) is mentioned in a recent publication referenced [below](#).

(Suggestions for improvement of any of my online documents/presentations are welcome!)

END OF BACKGROUND INFORMATION

Original Stepney-inspired diagram, followed by later versions

Below I'll present several diagrams combined with text, in an attempt to explain a complicated collection of ideas about evolution and development in egg-laying vertebrate species.

From late March 2022, I began to use the first diagram [below](#), inspired by Susan Stepney's sketch, as mentioned above.

I later switched to a more complex sequence of diagrams shown [further below](#), in the hope of bringing out some of the ideas more clearly.

I now provide both the simpler version giving a simplified overview, followed by more complex diagrams, giving more details. But I suspect all this can be done better! (Suggestions welcome.)

[The key idea developed during 2022 was that newly evolved complexity in the organism required new layers of structure in the in-egg hatching processes, which in turn required new layers of increasingly species-specific control mechanisms used during hatching.](#)

In more recently evolved species of vertebrate egg-laying animals, there are more layers of development, with later layers using species-specific meta-mechanisms also specified in the genome, i.e. mechanisms for using the later genetic information to control later stages of embryo assembly.

I am claiming that later stages of hatching in vertebrate eggs are controlled by more recently evolved meta-mechanisms for interpreting and using more recently evolved assembly instructions.

So gene expression (in the egg) includes parallel paths of specification of different parts of the new hatchling contents, and increasingly complex specifications of mechanisms for using increasingly complex, more recently evolved genetic details. The genetic specification controls not only the formation of the animals physiology within the egg, but also the behavioural abilities shown by the new animal shortly after hatching.

The increasing complexity of more recently evolved biochemical/biological mechanisms for gene-expression and control of gene-expression corresponds loosely with the increasing complexity of bootstrapping mechanisms in computer science and computer systems engineering since the 1950s, documented (incompletely!) by Wikipedia:

<https://en.wikipedia.org/wiki/Bootstrapping>

(Many computer scientists and computer systems engineers think only about [subsets](#) of these mechanisms.)

The original diagram below inspired by Susan Stepney's sketch showed evolution as going from left to right and in-egg development going from top to bottom. The later diagrams swap these: development goes from left to right and evolution from top to bottom. I think this made the diagram easier to understand (partly because English text flows from left to right!).

Use of virtual machinery

The presentation makes use of a notion of "virtual machine" that has emerged from a collection of increasingly complex developments since about 1960 that allow computational processes of many sorts, to be [implemented](#) in increasingly complex networks of physical machinery that can change over time while the virtual machines persist and change in different ways, often running on changing, possibly cross-continental, physical computing machinery.

Examples of such virtual machines include online banking systems, international email systems, flight reservation systems, trans-national air-traffic control systems, and, more recently, systems like Zoom, that make possible international, including trans-continental, meetings that bring together, at short notice, collections of users with different background knowledge, different hardware, in different locations, for lectures, discussions, family get-togethers, and many other purposes that don't require physical contact but make use of visual, auditory and textual communication across a very complex international network of information processing mechanisms and (constantly changing) information stores. For more information on Zoom see:

[https://en.wikipedia.org/wiki/Zoom_\(software\)](https://en.wikipedia.org/wiki/Zoom_(software)). [XXXXXX]

Beware of oversimplified/out-of-date notions of Virtual Machinery

Most of the online information about virtual machinery is out of date and over-simplified. A useful exception is this tutorial (which I have so far only sampled in part):

<https://www.vmware.com/topics/glossary/content/virtual-machine.html>

It describes several different kinds of virtualisation, but does not cover all the types that have become available during the 21st century, and of course, has no mention of the much older uses of virtualisation discovered millions of years ago by biological evolution.

The presentation below claims that long before human scientists and engineers discovered the powers of virtual machinery, biological evolution "discovered" and made use of the powers of virtual machinery in processes that transform the relatively homogeneous collections of matter inside an egg into the richly differentiated and enormously complex collection of interacting physiological structures in a developing organism, hatching in the egg.

I am not claiming that such mechanisms are used only in eggs -- there are many even more complex examples in other forms of reproduction, e.g. in mammals, but the key ideas are easier to present in relation to hatching processes in eggs of vertebrates because they achieve so much inside the eggshell, with so little interaction with the environment, unlike an embryo developing in a womb, for example, or gene expression in a plant seed during increasingly complex interaction with the environment as the plant grows.

This implies that, during individual development inside an egg, members of recently evolved vertebrate species, with multi-stage evolutionary histories, will go through several changes in control mechanisms for assembling various parts of a new organism during the hatching process.

As explained [below](#) this is not a variant of the theory according to which "ontogeny recapitulates phylogeny".

Why [virtual](#) machines?

Both

-- the control mechanisms used within the hatching egg for assembling physical particles to form the new hatchlings,

and

-- the control mechanisms for creating those new control mechanisms

must be *virtual machines* (in the sense evolved by 20th and 21st century computer systems engineering) because there is no room in an egg for machines composed of *physical matter* to manipulate millions of molecules to form the new physiological structures required by the developing embryo.

The virtual machines that achieve those results must somehow be *implemented in* physical/chemical machinery in the egg, without taking up extra physical space that new manipulating mechanisms would require. Such virtual machinery could be implemented using electromagnetic fields or signals, for example.

I suggest that there will be stages in construction of the new individual, corresponding to major transitions in their evolutionary history, and part of what happens at each stage is construction of new controlling virtual machines as well as construction of new chemical structures produced by those virtual machines as the physiology in the foetus becomes more complex.

This also raises the question: Which such chemistry-based virtual machines continue to play important roles in cognitive processes after hatching -- roles that apparently have never been noticed by neuroscientists? Please let me know if I am wrong. Roles of hormones in later stages of human development are special cases.

(Information about varieties of virtual machinery should by now be a standard part of education in philosophy, biology, psychology, neuroscience, social science, as well as computer systems engineering. But that has not happened.)

NOTE: Neither evolution nor development can be continuous

Evolution and development must both be at least partly discrete insofar as the mechanisms of change are chemical, and chemical changes include formation and removal of bonds between

physical particles: discrete changes.

The claim that reproductive processes that produce biological evolution must be at least partly discrete was clearly stated by Erwin Schroedinger in his 1944 book *What is life?* and the updated 1967 version:

https://archive.org/download/WhatIsLife_201708/What%20is%20Life_text.pdf

It is possible that the point was made earlier by one or more other scientists.

This implies that much that is written by philosophers of physics uses notations that cannot capture the mixture of continuous and discrete physical processes producing changes in complexity, used by mechanisms of biological evolution and development.

Evo-Devo diagrams

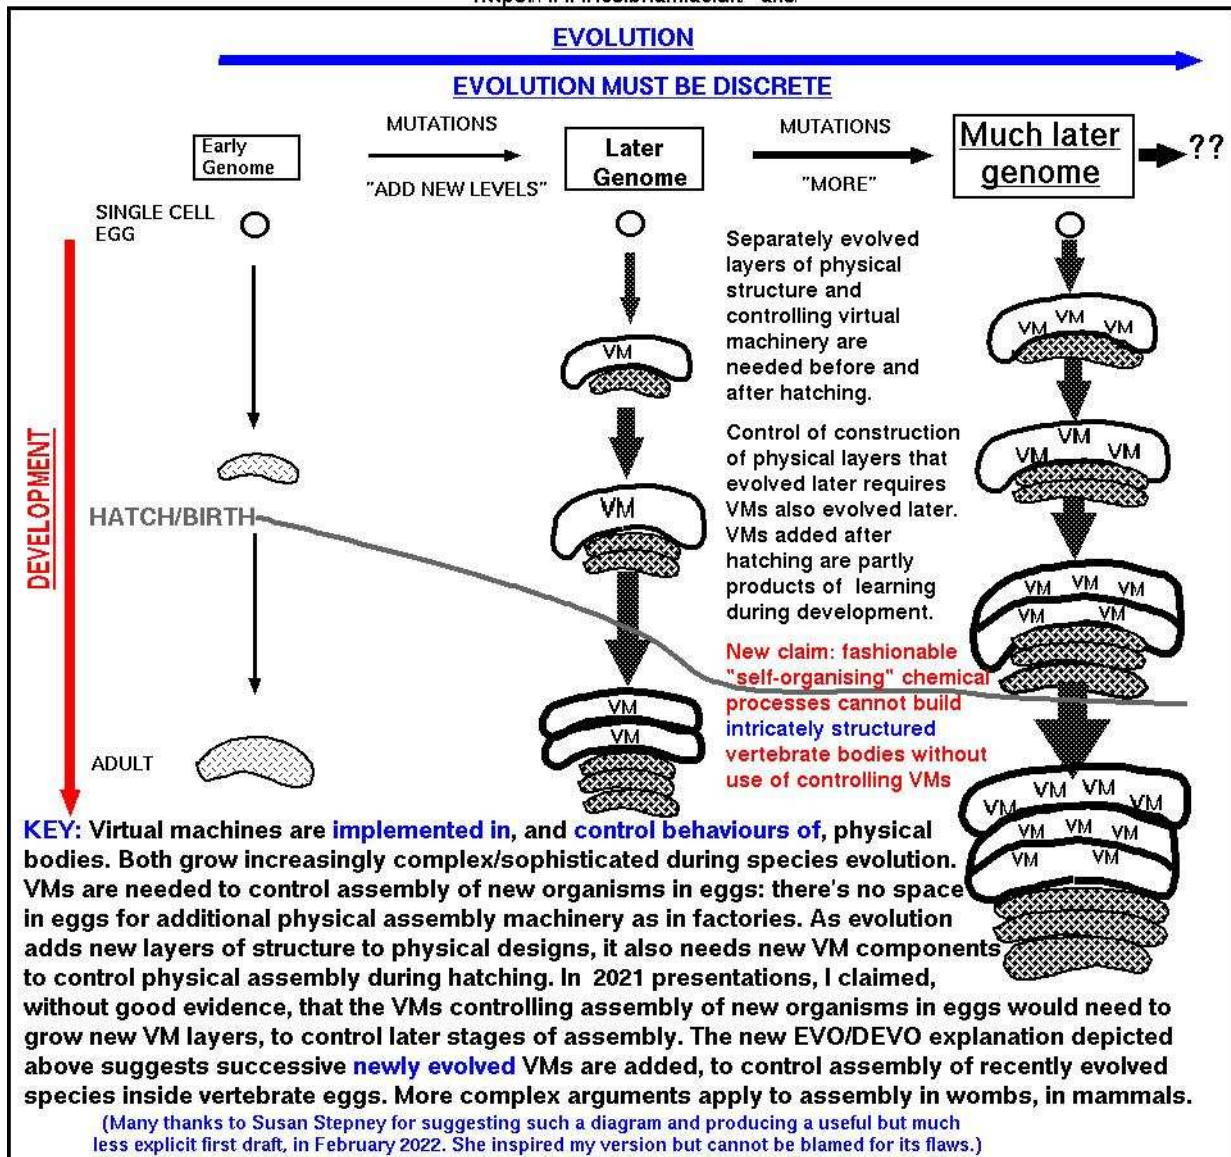
The first diagram below, inspired by Susan Stepney's pencil sketch (which had less detail), was my first attempt to represent the combination of evolutionary and developmental processes and mechanisms, using collections of virtual machines indicated as (VM) below.

Attempts to add more detail to the diagrams led to the revised format presented [further below](#).

Stepney-based Evo-Devo Figure

WHY DO THE IN-EGG ASSEMBLY PROCESSES IN EGG-LAYING VERTEBRATES REQUIRE A SUCCESSION OF DISTINCT (VIRTUAL MACHINE) MECHANISMS CONTROLLING ASSEMBLY?

Aaron Sloman, University of Birmingham, UK
<https://www.cs.bham.ac.uk/~axs/>



The figure above, inspired by Susan Stepney's depiction ideas, was used several times in talks during early 2022. Later I switched to the figures below. All are hard to take in without a verbal presentation.

A NEW COLLECTION OF EVO-DEVO DIAGRAMS

The figures below, produced in September 2022, are more recent, presenting a different view of the combinations of evolution and development.

[The third diagram below](#) labelled *Figure Evo-Devo (latest stage of evolution)*, summarises the changes in development across evolutionary transitions in a different way from the diagram above: [Figure Evo-Devo-Multi-VM](#).

At first I used only one new diagram, namely the third diagram but after deciding that it was too complex to be absorbed I added the first two diagrams showing how the third diagram is built up. Readers who find the third diagram below intelligible can skip the previous two!

The first diagram below summarises ideas about an early stage of evolution of an ancestor of current vertebrate egg-laying species, and adds some hints about effects of later evolutionary changes produced by genetic mutations.

The second diagram shows a later stage of evolution, in which there is a major [new transition during development](#) of the foetus inside an egg, requiring a transition to [a new, more complex, form of control of development](#) because of the need to coordinate development of more, and more diverse, components.

[The third diagram below](#) subsumes the two preceding diagrams, and adds a much later stage of evolution, where processes controlling developments in the egg have several new layers of complexity, because the physiological structures being assembled are more complex and diverse -- as a result of more evolutionary transitions.

If a better form of presentation occurs to me, or the ideas evolve, I may change these diagrams later. (Suggestions welcome!)

Figure Evo-Devo-a (early evolution)

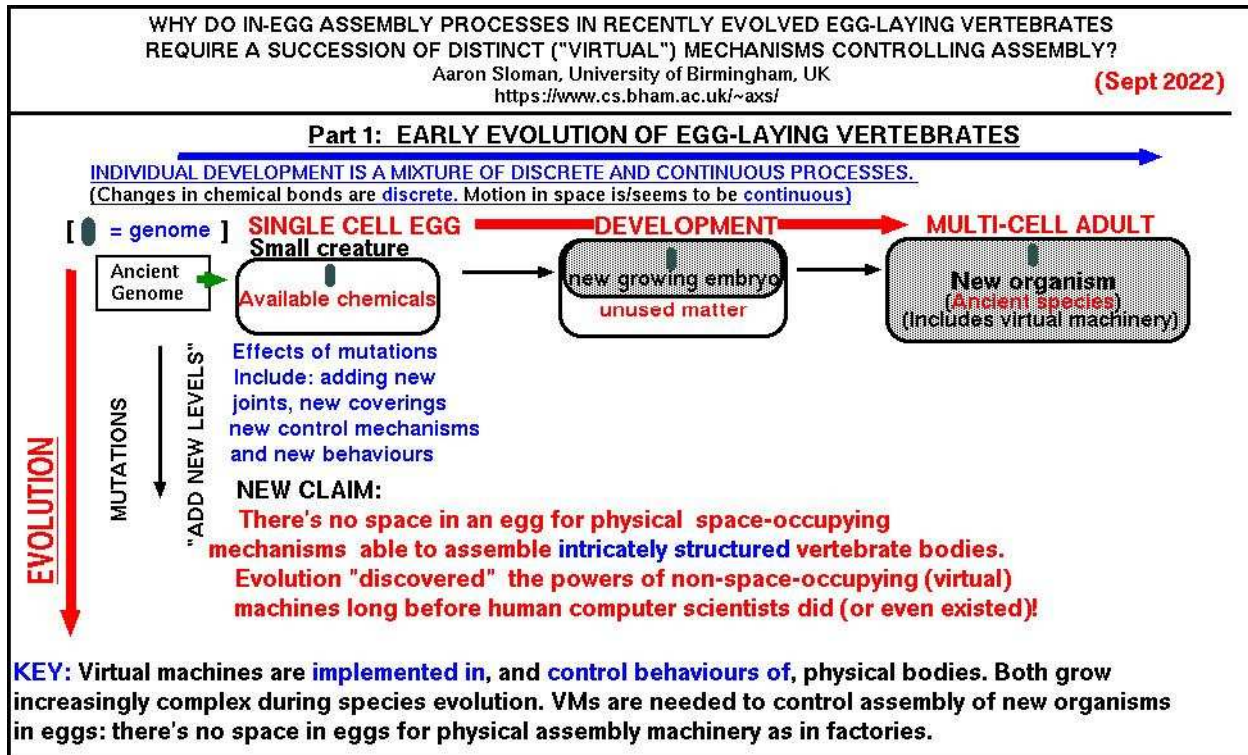


Figure Evo-Devo-b (later evolution)

More than one stage of development.

Note: this figure and the next one slightly modified 16 Sep 2022

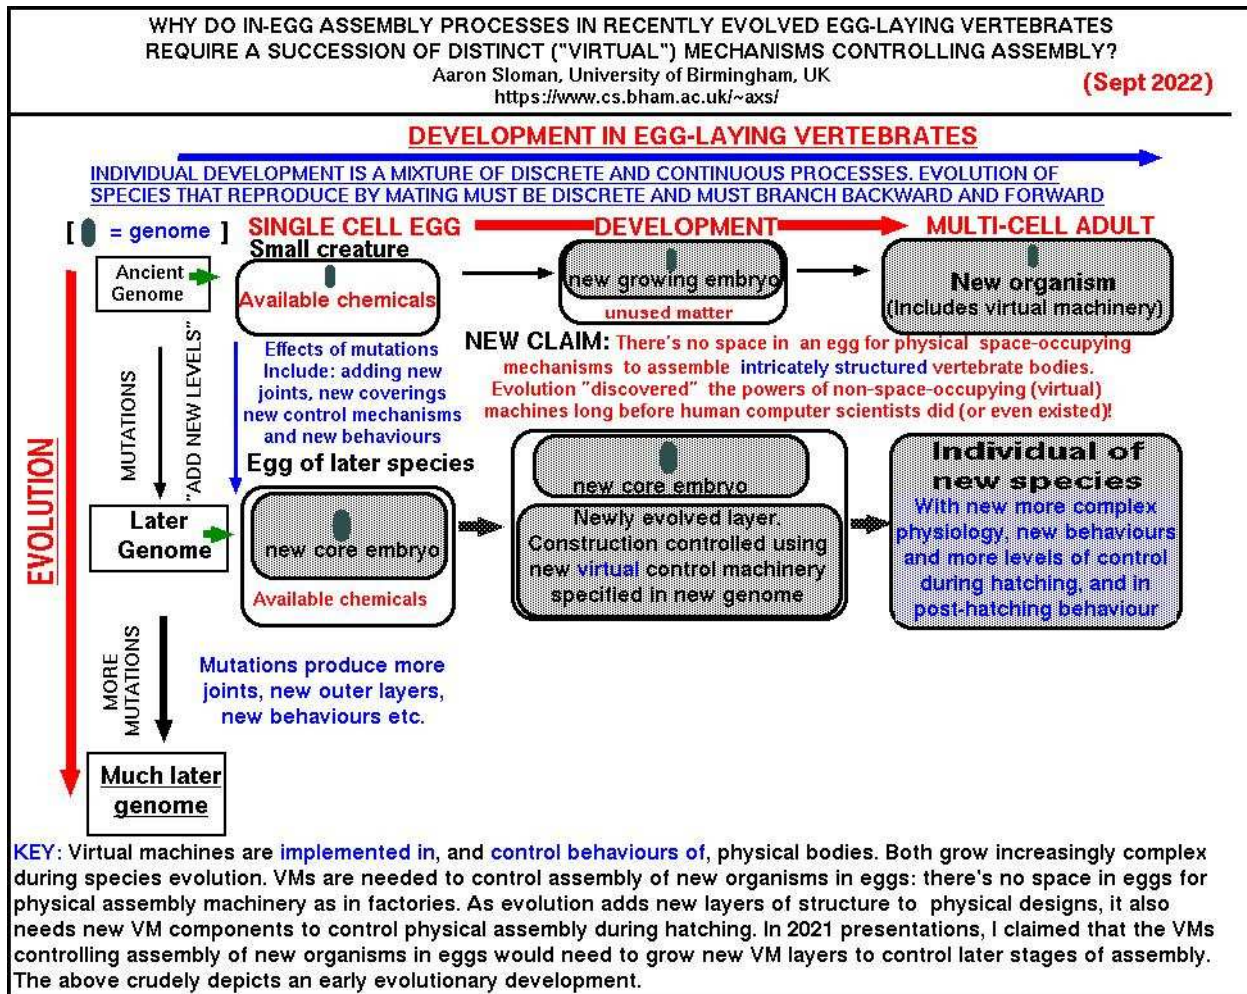
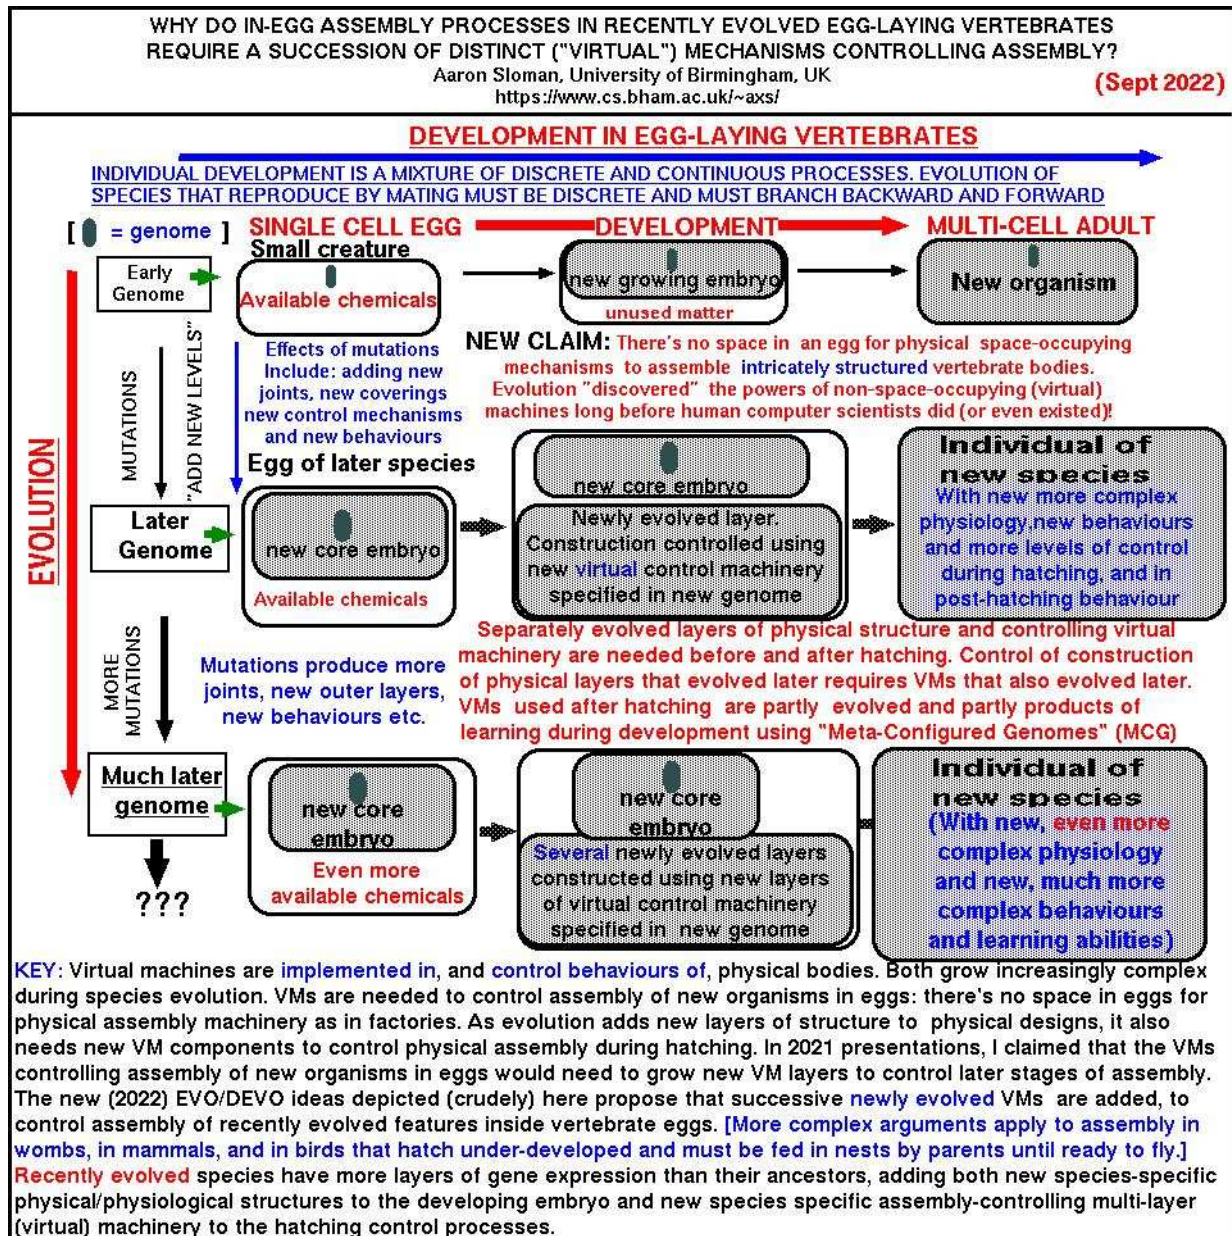


Figure Evo-Devo (latest stage of evolution)

Including more layers of evolution, and corresponding new layers of development, combining stages depicted above.



Changing Control Ontologies

A feature of the theory that is not captured in the above diagrams is that the mechanisms involved in controlling developmental processes in the eggs must use different kinds of information, including:

- information about the current state of development in various parts of the embryo,
- information about states that do not yet exist but need to be brought into existence, and
- information about possible changes that could be made to occur, or should be prevented from occurring, i.e. information about goals and preferences: **control information**.

As the reproductive process continues, new kinds of information become available for use in controlling embryo-construction actions. Such useful information includes information about what has so far been achieved, derived from internal sensing, i.e. [factual information](#) and information about what needs to be achieved in the current situation, derived from species-specific control information in the genome, which only becomes accessible to action-selection processes at later stages of reproduction.

Such factual and control information can be used to evaluate current states of affairs, e.g. as desirable, incomplete, undesirable, or better or worse than some other possible states. Such evaluations can then be combined with factual information about possible actions in the current situation, or target information, about future states to be achieved. This can drive control processes that select changes to be made, or changes to be prevented (e.g. which states should be preserved, once achieved).

In abnormal cases, where something unusual has happened during development, unusual compensatory developmental processes may be triggered. (How? Different cases need to be distinguished and investigated.)

Such varieties of information about what has been achieved so far and about what needs to be achieved, or prevented, will depend on the current state of development of the embryo, and will keep changing over time.

As the embryo grows, new controlling sub-processes will be generated, combining information about what has been developed so far with genetic information about what to add, which will be different at different stages of development in an individual and will also differ across species.

A new control layer will need to use information about what materials and structures have already been assembled in different locations and which new ones need to be assembled in different locations, and how the components of the new ones need to be related to old components (e.g. extending, or branching, or attaching a muscle, or providing new blood vessels or nerve fibres for the new anatomical structures).

[As the parts of the new embryo become increasingly differentiated the variety of types of information about what exists and what needs to exist will also become increasingly differentiated, as will processes controlling changes in different parts of the developing embryo.](#)

That may (occasionally?) include information about something that exists that should not exist and therefore needs to be undone or compensated for in some way.

For example, if a developing embryo for some reason starts to develop two heads, this will require changes in control of other developments, so that the two heads become connected with appropriate resources, including blood supply, nerve fibres, bone structures, muscles, etc. In some cases the compensatory changes do not succeed and the foetus dies, whereas in others a complex collection of coordinated adjustments to standard developmental processes occurs and a highly abnormal individual emerges, which may or may not survive for some time after hatching.

Is all this a version of recapitulation theory?

(Added 17 Sep 2022)

The ideas presented above may appear to constitute a variant of the discredited "Recapitulation Theory", which claims (roughly) that the development of individual organisms recapitulates the evolution of their species, summarised as "ontogeny recapitulates phylogeny". Wikipedia explains the idea: https://en.wikipedia.org/wiki/Recapitulation_theory.

I am not claiming that ontogeny recapitulates phylogeny. The relationship between developmental processes and evolution is far more complex and more subtle than mere "replication".

What I have been claiming is that *for many vertebrate species that lay eggs, development of the mechanisms controlling development of an individual organism partly replicate the mechanisms and processes used during the evolution of the species*. This could perhaps be called "[Process recapitulation theory](#)", in contrast with "[Result recapitulation theory](#)".

One of the implications is that the earliest stages of gene expression in these organisms use mechanisms that have a lot more in common across species than the mechanisms used in later stages. This is most obvious for species that emerge from their shells with very different outer coverings e.g. feathers, shells, scales, etc. and very different behaviours (e.g. walking, crawling, swimming).

The very earliest stages of replication, i.e. splitting and duplication of the initial DNA in one cell of the new-laid egg have been studied in great detail and have much in common across different species that use sexual reproduction. But as reproduction within an individual egg continues, the process-patterns that unfold in different species, and the mechanisms used to achieve those developments, diverge more and more.

The aim of this document is to draw attention to that divergence, its origins, and its implications, while admitting that although the processes diverge widely across species, there are some higher order common patterns in that divergence, shown clearly (I suspect) in the reproductive processes of egg-laying vertebrate species that hatch with fully functional bodies and a collection of useful cognitive competences that don't come from learning in the environment.

Perhaps my claims could be given a label something like "[Meta-process recapitulation theory](#)".

Associated with those general goals is an attempt to draw attention to the depth and the powers of species-specific mechanisms of reproduction that produce not only enormously complex, species-specific, [physiological structures](#) in newly hatched individuals, and species-specific [patterns of physiological development after hatching](#), but also produce species-specific [cognitive competences](#) available shortly after hatching without requiring learning or training, including perceptual competences, goal formation competences and goal achievement competences.

Example: video of newly hatched avocets

It is not clear to me whether current physical theory can explain how those in-egg developments are controlled, especially insofar as the in-egg processes somehow produce not only complex physiological structures but also complex [post-hatching competences](#) that do not need to be learnt, e.g. the abilities of newly hatched avocets shown in this 35 second videoclip from the BBC Springwatch programme in June 2021:

<https://www.cs.bham.ac.uk/research/projects/cogaff/movies/avocets/avocet-hatchlings.mp4>

How can chemical assembly processes in an egg produce an avocet that has not only enormously complex and intricately interconnected internal physiological structures and mechanisms, but also motion and feeding competences available for use without having to be learnt?

I am constantly amazed by the number and variety of researchers who mistakenly believe that such competences in animals have to be acquired by training neural networks to derive consequences from sensory and motor data. Such theories, proposed or believed by many distinguished researchers, are clearly refuted by the competences of many newly hatched individuals of many species. They don't have to learn how to move around, detect food, eat it, etc.

There are clearly mechanisms that are much more powerful than trainable neural networks and are used to produce a huge variety of post-hatching competences in the young of many different species, with different physiologies, shapes, sizes, external coverings, different environments, different forms of behaviour and different requirements for food.

There are also many differences in post-hatching competences of different egg-laying species, just as there is a huge variety of physical forms, including many types of bird, alligators, turtles, snakes, ... and many more.

At present I know of no attempts to describe developmental mechanisms that could explain how all those competences are produced by mechanisms in eggs. Or how related competences could be produced in wombs of mammals, e.g. new-born foals able to run with the herd within a few hours of birth!

These explanatory problems appear not to have been noticed by most philosophers and scientists (e.g. psychologists and neuroscientists) working on cognition. Please inform me if you have encountered or produced work that states the problem (explaining spatial cognitive competences of newly hatched or newborn animals) and provides, or attempts to provide an explanation.

Note:

There are also egg-laying species whose newly hatched young are NOT ready to move around and find food, e.g. birds that hatch in nests in trees or on cliff faces. For them, moving around successfully involves flying, and that requires development of powerful muscles that are not needed for newly-hatched walkers or crawlers. So their post-hatching behaviours are restricted to acquiring and consuming food brought to the nest by parents, until they have strong enough muscles to support flying.

IMPLICATIONS FOR FUNDAMENTAL PHYSICAL THEORY???

Are the assembly mechanisms used during different stages of hatching, comparable to increasingly sophisticated Maxwell Demons?

Thinking about the processes mentioned above led me to wonder whether features of hatching processes may have previously unnoticed implications for fundamental physical theory, although I am still groping for a good way to characterise those implications.

One of the key thoughts is that the mechanisms that control the (enormously complex) biochemical changes involved in controlling development of the new organism inside the egg are partly similar to the much simpler mechanism described as "Maxwell's Demon" in a thought experiment attributed to James Clerk Maxwell, summarised here:

https://en.wikipedia.org/wiki/Maxwell's_demon.

However, the in-egg demons, especially the later versions, are far more complex: more like a self-extending sequence of Maxwell-demons, each of which produces some new order inside the hatching egg and starts the process of creating a new, more complex, more recently evolved, demon, or collection of demons, required to control later processes of assembly of physiological structures and mechanisms in the developing egg. The diagrams above crudely represent such layered processes of gene expression.

Types of development that need to be controlled at later stages include growth of bones, of muscles, of tendons, of nerves, of blood vessels, of skin or other covering material, etc. Moreover, these developments are all parts of larger developments, each of which forms a subsystem that requires coordination of its own components, as well as coordination between systems. Examples of such sub-systems could include:

- outer covering materials (skin, hair, fur, feathers, scales, shells, etc.),
- blood transport and delivery systems, including arteries, veins, capillary networks where veins and arteries merge, and the blood pumping system, and after hatching transportation systems for the oxygen brought in via lungs and carbon-dioxide exhaled via lungs,
- networks of nerve-fibres carrying information signals in various directions, to control and coordinate internal and external behaviours.
- others ...

Many philosophers who write about physics assume that the mathematical structures required for representing the physical structures and processes in organisms will need to represent relationships between numerical values and rates of change of numerical values. But such forms of mathematics are not enough: there is also a need to represent assembly and disassembly of structures of various sizes and also relations between structures, including contiguity, containment, attachment, and local transfer of forces.

Clearly a level-N demon cannot already have the knowledge required for the level-N+1 assembly task. The information is presumably available in portions of the DNA that have not yet been used to control assembly. But a new more sophisticated, more recently evolved control process for the level-N+1 task has to be triggered by the level-N demon: it "knows" how to create the level-N demon that will have the competences required to control the level-N+1 assembly processes.

Is the comparison with Maxwell's demon mistaken because the mechanisms controlling increasingly complex forms of differentiation in the egg can use chemical energy liberated when complex molecules provided in the new-laid egg are decomposed to provide the fragments used to construct new components of the physiological structures in the developing embryo?

A question about post-hatching results

Is it possible that the mechanisms outlined above help to answer the question: What makes it possible for many newly hatched animals to have both fully formed bodies **and also** cognitive functions that enable them to behave appropriately in the environment, without having to train neural networks or undergo any other form of learning, like the newly hatched avocets mentioned [above](#).

I suspect the answer to that question will turn out to be that there is an additional developmental process that evolved as a side-effect of the mechanisms outlined above. Perhaps the evolutionary processes that enabled in-egg assembly processes to be controlled in order to produce a duck, chicken, alligator, turtle, etc. were copied and modified in a manner that allowed them to support post-hatching manipulations of **external** physical objects instead of only **internal** manipulations of components of physical structures assembled during production of a new animal in an egg.

Turning that very crude idea into a precise theory that can be tested is a major challenge for this project.

DRAFT ADDITIONS

The remainder of this document is still an early draft. I hope to be able to include references to more related work in future. Suggestions welcome.

Background 1: The Kant/Hume Disagreement

One of the key disputes about the nature of mind or cognition was Immanuel Kant's disagreement with David Hume's claims about the contrasts between

-- kinds of knowledge that are either (a) about "matters of fact" or (b) about definitional "relations between ideas",

and

-- supposed advances that are neither of type (a) nor type (b), which Hume claimed to be "mere sophistry and illusion", including theological arguments.

The contrast between (a) and (b) is sometimes referred to as "Hume's fork".

David Hume and Immanuel Kant (from Wikimedia)



Kant criticised Hume by claiming that there are significant discoveries, including ancient mathematical discoveries, that are in neither branch of Hume's "fork". Those discoveries are not based solely on empirical observation and they are not discovered merely by deriving logical consequences of (explicit or implicit) definitions of the concepts used combined with purely logical assumptions. Such knowledge is therefore neither empirical nor analytic.

Moreover those discoveries are not about contingent truths or falsehoods but about **necessary** truth or falsehood.

Examples include Pythagoras' theorem and other discoveries in geometry and topology made centuries before well known ancient mathematicians such as Pythagoras were born. Another example is the discovery that one-to-one correspondence is a transitive relation between sets: not understood by humans until they are five or six years old (as Jean Piaget discovered).

Such discoveries are

- synthetic, not analytic, i.e. not simply based on logic and definitions,
- non-empirical (i.e. *a priori*), and
- necessarily true (i.e. they are non-contingent: counter-examples cannot exist).

Kant seemed to think it was impossible for humans to understand the mechanisms making such discoveries possible. He suggested that the mechanisms would lie "forever concealed in the depths of the human soul".

I suspect that if he had lived two centuries longer he might have proposed reasoning mechanisms supported by brain chemistry -- not the currently fashionable "neural network" mechanisms that merely collect statistical evidence and then derive probabilities from the data. Such neural networks are constitutionally incapable of producing proofs of necessity or impossibility. They are restricted to discovering low or high probabilities.

The current majority view among philosophers (especially those lacking a good education in geometry and topology) seems to be that Kant was mistaken in claiming that we can discover necessary truths that are neither empirical nor simply logical consequences of definitions.

I think Kant was not mistaken, but a full defence of his views will require progress in the Meta-Morphogenesis project, investigating chemistry-based reasoning mechanisms, on which I suspect Alan Turing had been working shortly before he died. I am not referring to Turing's work on chemistry-based 2D pattern formation published in 1952 which I think was simply a sideline in a deeper investigation. Those ideas are discussed in this long, messy document, which will later be revised: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/hatching-talks-2022.html>

That work provides some of the long term motivation for the ideas presented here, concerning processes of development in eggs of vertebrate species whose young have complex competences available soon after hatching, which they could not have learnt after hatching. The competences must be products of hatching processes in eggs.

Background 2: The Meta-Morphogenesis project

The ideas presented here are closely related to, and emerged from, the *Meta-Morphogenesis* project, inspired by some of Turing's ideas. The project was triggered around 2011 when I was asked by Barry Cooper, the main editor of a [commemorative volume](#) on Alan Turing, to comment on Turing's 1952 paper on 2-D chemistry-based morphogenesis (i.e. 2D pattern formation on the surfaces of organisms).

My commentary conjectured that Turing's study of processes producing changing 2-D surface patterns turned up merely as a "side issue" for Turing, during much deeper, more complex, still unpublished, research he was doing at the time, on chemical mechanisms involved in controlling reproduction and development of biological organisms, closely related to conjectured chemical mechanisms underpinning ancient human mathematical reasoning about spatial structures and processes.

These ideas are also relevant to spatial intelligence in many other species and are closely related to earlier work with Jackie Chappell on "Meta-Configured" genomes, referenced [below](#).

One of the aims of all this work is to show that current theories of intelligence based on neural networks (NNs) are seriously misguided, partly because many newly hatched animals (like the avocets shown in the BBC videoclip [below](#)) display important competences that they have not had time to acquire by training neural networks, and partly because some of those competences involve detecting impossibility or necessity, which *cannot* be determined on the basis of statistical evidence or derived probabilities: impossibility and necessity are not extremes on a probability scale. They cannot be detected by neural net based mechanisms that collect statistical evidence and compute probabilities.

This seems not to be understood by most researchers who investigate learning and reasoning mechanisms based on trainable NNs. Many of them have never studied geometry and topology using diagrammatic proofs because of disastrous changes in mathematical education around the middle of the 20th century.

The most popular alternative to both diagrammatic reasoning and empirical learning in mathematics is use of logic-based symbolic reasoning. This has led to many deep mathematical results and development of computer-based mathematical reasoning systems.

But that still leaves unexplained the mechanisms used in ancient forms of mathematical discovery using spatial reasoning.

I think we can learn about a potential role for chemistry-based mechanisms in providing explanations, if we study hatching mechanisms in eggs of vertebrates, as sketched below. I did not recognize the significance of hatching processes in eggs until around September 2020.

In fact much of the relevant evidence is already widely known, but its significance goes unrecognized.

Moreover the mechanisms that make such mathematical spatial reasoning possible are not yet known. Statistical evidence may persuade a neural network that something is false, e.g. it is false that a planar polygon can have more vertices than edges. However statistical evidence cannot establish impossibility.

What sort of mechanism could enable a neural network to establish that some proposition is **necessarily** false, i.e. its truth is impossible. Statistical evidence is irrelevant to proof of necessity or impossibility, though it can be relevant to proof of desirability. So we shall have to explore alternative forms of evidence.

This Document

This document is available at:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evo-devo-figs.html>

Since September 2020, I have been investigating, and giving talks about, processes and mechanisms that I conjecture are involved in reproduction using eggs in many vertebrate species. That was when I first realised that hatching mechanisms were relevant to some deep unsolved problems about spatial intelligence. As far as I can tell, nobody else has noticed the connection with hatching.

[If I there are others, please send me details (a.sloman AT bham.ac.uk).]

The restriction to vertebrate species does not imply that the ideas are not relevant to eggs of non-vertebrates, e.g. insect eggs, or plant seeds. I focus on egg-laying vertebrates because that includes so many different species, with widely varying physiological forms and complex, but different, post-hatching behavioural competences, e.g. species-specific abilities to control spatial movements and feeding processes shortly after hatching. That indicates both the wide applicability and the many detailed variations in the mechanisms used.

Increasingly complex extensions to those ideas about hatching processes, were presented in "evo-devo" talks during 2021.

In January 2022 a new strand emerged, investigating changing relationships between [evolution of egg-laying species](#), and [chemical developments during hatching of each egg](#).

It was not clear to me whether current physical theory could explain how those in-egg developments were controlled, especially insofar as the in-egg processes somehow produce very complex and intricately interrelated physiological structures [and also produce complex post-hatching abilities that do not need to be learnt, including the motor abilities of newly hatched animals of many kinds, such as chickens, ducks, avocets \(illustrated in the videoclip below\)](#), turtles and many other species of egg-laying vertebrates.

How can chemical assembly processes in an egg produce an animal that has such competences -- competences that many researches nowadays (mistakenly) believe have to be acquired by training neural networks to derive consequences from sensory and motor data? Those beliefs, held by many distinguished researchers, are clearly refuted by the competences of many newly hatched species.

A useful, but shallow, introduction to hatching processes in eggs is provided by this 2013 video showing chicken embryo development: <https://www.youtube.com/watch?v=PedajVADLGw>

Here's a video from China showing a chick hatched without an eggshell, in a transparent bag, in 2018: <https://www.youtube.com/watch?v=j0XmhPZwMuA>

This collection of images shows stages in the development of a chick in an egg:

<https://i.pinimg.com/736x/62/28/a6/6228a61f6607dcd92fcc775ccf8cd867--chicken-life-a-chicken.jpg>

It is impossible for a few such videos and images to convey the huge variety of patterns of development that occur in eggs of different vertebrate species, including variations in body form, in outer coverings, and in post-hatching behaviours, e.g. in turtles, alligators, various kinds of snakes, chickens, and the avocets shown in a video-clip [below](#).

I have not encountered anyone else (apart from Alan Turing if my guess above is correct) who takes seriously the suggestion that such post-hatching competences are produced by chemical assembly processes in eggs, though I suspect Immanuel Kant would support this suggestion if he were still alive. The relevant biochemical knowledge and recording technologies did not exist when he was writing.

My thanks to Professor Susan Stepney (York University)

<https://www.cs.york.ac.uk/people/susan>

As mentioned above Professor Susan Stepney (York University), with whom I have discussed related ideas over several years, after hearing one of my talks early in 2022, sent me a pencil drawing attempting to summarise my latest ideas in a 2-D array of evolutionary and developmental changes, with evolutionary changes shown horizontally from left to right and developmental changes shown vertically from top to bottom. She developmental processes had changed during evolution -- and helped her to understand the claims I was making!

That inspired me to incorporate a diagram, inspired by her sketch, in presentations from June 2022.

Over several months, across several presentations, I made attempts to improve the theory and the diagrams, by adding more detail and attempting (and perhaps failing?) to make an increasingly complex collection of ideas easier to understand.

There is a lot of research on related but much simpler problems

There are now many researchers investigating organisms that are capable of re-organising their physical structures (e.g. slime moulds), but I have not encountered researchers who attempt to explain not only how self-organising hatching processes in eggs can produce extremely complex physiological structures but also how they provide newly hatched individuals with important [behavioural and cognitive competences](#), like the newly hatched avocets mentioned above, and many other species that display complex forms of (species-specific) spatial intelligence shortly after hatching, used in activities such as walking, crawling, feeding, following the mother, etc.

E.g. the otherwise excellent work of Mike Levin [below](#) does not address, or even mention, this problem of explaining complex emergent spatial competences, as far as I can tell. And none of the self-assembly processes studied by other researchers whose work I have encountered (e.g. processes in slime-moulds) involve the degree of [intricate structural differentiation and re-organisation](#) produced in eggs during hatching.

Another relevant thread: Meta-Configured Genomes

A collaboration with Jackie Chappell

My work on hatching mechanisms began in 2020 but relates to a much older investigation of "Meta-Configured Genomes" developed in collaboration with Dr. Jackie Chappell (<https://www.birmingham.ac.uk/staff/profiles/biosciences/chappell-jackie.aspx>) after she came from the Ecology department in Oxford to Biosciences at the University of Birmingham in 2004. Some of the ideas we developed after she arrived were published in [Chappell and Sloman\(2007\)](#).

Some of the later threads connected with biochemical details were suggested by Prof Peter Tino (<https://www.cs.bham.ac.uk/~pxt/>) around 2019.

The collaboration with Jackie led to the idea of a [Meta-Configured Genome](#), containing aspects of gene-expression that provide patterns or templates expressed at various stages increasingly late in life. We rejected Waddington's idea of development as motion along a fixed "Epigenetic" landscape specified by the genome: it failed to capture the fact that the same genome could have very different products in different environments, as illustrated by the diversity of languages used by humans who presumably share a mostly common human genome, including spoken languages, sign languages, written languages and languages created for special purposes, e.g. mathematics, computer programming, scientific theories, etc. We used the label "Meta-configured genome" (MCG) to draw attention to the genetic bases of these capabilities.

A video tutorial presenting some aspects of these ideas is available here:

<https://www.cs.bham.ac.uk/research/projects/cogaff/movies/meta-config/metaconfig.webm>

MCGs are not [complete](#) genetic specifications, because they don't contain all the information required for gene-expression. Instead some of the genetic material expressed at a relatively late stage of development contains "gaps" that are filled using information acquired at earlier stages during individual development while the individual was interacting with the [then current](#) environment, i.e. [not](#) the environment that existed when the genes first evolved.

So expression of a 6 year old's meta-configured genes will produce "gaps" that are filled using information picked up earlier by that individual --- information that can vary across geographical locations and across species history in a fixed location. For example, genetic mechanisms controlling learning during play activities of a 6 year old can use information acquired earlier (e.g. during play, or social interaction) by that individual in that location, and the results can differ widely across geographical locations and across human history in a fixed location.

Some of the differences between cognitive development of humans born during the 21st Century and cognitive development of their parents (who could not have encountered mobile phones, email, or the internet) illustrate this point.

Differences between languages, including sign-languages used in deaf communities and 'click' based languages in Southern Africa also illustrate many effects of meta-configured genomes. (See https://en.wikipedia.org/wiki/Khoisan_languages)

The use of a meta-configured genome allows later processes of gene expression to be partly genetically determined by evolutionary changes many centuries earlier, and partly tailored to details of the current environment, including details recorded by the learner during earlier stages of individual development. Those details may depend on recent history of changes in the learner's environment. Such an environment can fill gaps in (provide parameters for) portions of the genome that are expressed later and provide generic capabilities that are partly instantiated using

information picked up by the individual, e.g. from the physical and/or social-linguistic environment.

This process can have several developmental layers, so that individual development includes several stages at which meta-configured genomes are activated, whose interaction with the environment provides information used later, during later expressions of (more abstract) meta-configured genome features.

Perhaps the most obvious and spectacular example of this is the way in which genetically specified features of human language development that occur relatively late in life can vary enormously across individuals that are born at different times or in different locations. (Or both!)

In particular, later stages of language development at any time can depend on how features of language development in the current culture have changed in the past. This explains how young humans with the same language-related genetic mechanisms can develop enormously varied detailed [linguistic](#) competences in different location on the planet, and also how, across generations in the same culture, a language can change because each new generation picks up, makes use of and possibly extends, modes of linguistic communication and reasoning that were not available to their distant ancestors.

This also helps to explain how use of language by humans includes signed languages developed in deaf communities (or sub-communities) and written languages, as well as accounting for the huge diversity of spoken languages that have developed among spatially separated groups of humans.

(I have argued elsewhere that, in human evolution, sign languages used in cooperative activities must have developed before spoken languages. Other phenomena to be explained include human abilities to create entirely new spoken or written or printed languages for particular applications, e.g. in mathematics, science or computer programs.)

We (Jackie Chappell and I) claim that these ideas apply not only to human linguistic competences, but also other competences that are built up in stages in ways that allow previously evolved patterns to be instantiated in ways that are profoundly influenced by developments in earlier generations that altered the developmental environment for later generations --- as in the history of technology, science and mathematics.

This is totally different from theories that assume that some uniform learning mechanism starts from scratch with each newborn individual, collecting data, and deriving statistical patterns, as in most current "Neural Network" models of learning.

Many researchers appear not to have noticed that such mechanisms can discover only statistical regularities with associated probabilities. They cannot explain ancient and not so ancient mathematical discoveries concerning mathematical [impossibilities](#) and mathematical [necessities](#), for example, discoveries in geometry and topology, such as proofs of Pythagoras' theorem, or discovering that spatial containment is necessarily transitive. Many such mathematical discoveries were made centuries before Pythagoras was born! Those ideas were almost certainly related to discoveries about necessity and impossibility that were first made in contexts of performance of more or less complex physical tasks, for example the tasks of building large constructions using materials that had to be transported long distances cut into appropriately sized and shaped components and then moved into final locations on partly constructed buildings, including large pyramids, temples, bridges(?), aqueducts(?), etc.

Related forms of relatively recently developed human activities using new languages or formalisms tailored to specific mathematical, scientific or engineering tasks, using linguistic structures that do not occur in any previously developed human languages include modern programming formalisms of many different kinds mainly used not for communication between humans but for creation of new kinds of computer-based technology.

As far as I know, there are no theories in linguistics, philosophy, psychology or neuroscience, that explain how humans are able to create and use such languages and to create new forms of machinery that use them in processes with speed and complexity that cannot be matched by human thought processes. I.e. biological evolution has provided humans with abilities to create new forms of non-biological language users that can perform information-processing tasks that no humans can perform, individually or collectively.

Work to be done

Work for the future includes: combining the ideas about mathematical cognition with more detailed versions of the above proposals for mechanisms of gene-expression in eggs, and then using these ideas to give a new account of the mechanisms underlying ancient forms of mathematical competence and mathematical consciousness, which are not explainable

- either on the basis of statistics-based neural networks (which are constitutionally incapable of discovering cases of necessity or impossibility)
- or on the basis of modern formal reasoning mechanisms that are 20th century extensions of 19th century formal logic.

Those modes of reasoning are recent inventions and I know of no evidence that they were used consciously or subconsciously by the ancient mathematicians who first discovered theorems in geometry or the ancient engineers who used informal versions in their practical activities (e.g. transporting very heavy objects and constructing pyramids, temples, etc.).

Earlier work leading up to this

The above significantly extends the ideas presented in my 2020 paper Varieties Of Evolved Forms Of Consciousness, Including Mathematical Consciousness, in *Entropy* Vol 22(6:615), <https://doi.org/10.3390/e22060615>

That paper extended work on the Meta-Morphogenesis project, which was triggered around 2011 by an invitation to contribute a comment on Alan Turing's 1952 morphogenesis paper for the Elsevier volume celebrating the 2012 Turing Centenary. That book was published in 2013 <http://www.cs.bham.ac.uk/~axs/amtbook>.

My contribution, in pages 849-856, speculated that Turing's paper on 2D pattern formation merely reported a *side-issue* that had turned up in a deeper, unpublished, investigation of relationships between chemistry and brain functions. I later tried to specify that deeper investigation, labelled a study of "Meta-Morphogenesis" here:

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

(Also PDF)

Some online geometry/topology tutorials

-- Presentation on Euclidean geometry by Zsuzsanna Dancso at MSRI.

"Trisecting angles and calculating cube roots was a big problem for Euclid and his cohorts."

<https://www.youtube.com/watch?v=6Lm9EHhbJAY>

Now at University of Sydney

<https://www.msri.org/people/12337>

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

-- Topology & Geometry - LECTURE 01 Part 01/02 - by Dr Tadashi Tokieda
Including cutting a mobius strip down the middle: count the number of twists.
(He has many more online tutorials and demonstrations.)

-- Online lectures by Prof Cem Tezer (recently deceased)

Middle East Technical University

<https://www.youtube.com/watch?v=AJvjtK2mmpU>

MATH 373 - Geometry I - Lecture 1

<https://www.youtube.com/watch?v=1hNR-iCuw7g>

MATH 373 - Geometry I - Lecture 2

[There are many more online geometry tutorials.]

MORE RELATED WORK

(This is a tiny subset! Email a.sloman@bham.ac.uk with suggestions for inclusion in a more comprehensive list.)

A lot of relevant research has been done and is being done on self-organising biochemical systems. A small subset is referenced below. However, I have so far not found any such work that explains, or even tries to explain, how hatching processes in vertebrate eggs can produce new hatchlings that have extremely complex interconnected fully functional physiological components and have competences that don't need to be learnt in the environment, e.g. by training neural networks. For example they can move around, avoiding obstacles, following the mother (in some cases), and feeding themselves, like the newly hatched chicks and ducklings detecting and eating food while walking around and avocet hatchlings shown in the above [video clip](#) feeding in water.

Most of the detailed post hatching competences are species-specific so there must be a lot of cross-species variation in the transformations of egg contents during hatching that are relevant to cognitive competences, in addition to species specific hatching processes that produce variations in size, physical forms and modes of locomotion evident shortly after hatching in chickens, swans, alligators, crocodiles, turtles, various kinds of snakes, etc.

Proposed "self-organisation" mechanisms (e.g. training of neural networks) that I have encountered so far do not explain how the biochemical processes in eggs of vertebrates can assemble animals containing a large variety of different types of highly intricate, richly differentiated, functionally and spatially related, partly species-specific, physiological structures that exist in newly hatched vertebrates, in addition to sophisticated behavioural competences. (Compare [Hermann Rahn, et.al. \(1979\).](#))

I have also not encountered any recent attempt to explain the phenomena mentioned by Immanuel Kant, such as the human ability to use spatial reasoning about what is **impossible** or **necessarily the case**, as in ancient mathematical discoveries in geometry and topology.

NOTE: This web site focuses on vertebrate egg-laying species, but there are also important things to be said about processes in other sorts of eggs, e.g. insect eggs, cocoons, germinating seeds, etc. In principle it would be desirable to produce a survey of attempts to explain how hatching mechanisms work across a wide range of species. As far as I can tell most or all of the explanations produced so far regarding post-hatching physical forms and behaviours refer only to the early stages of gene-expression that are common to a wide variety of species, before highly species-specific structures are produced in the eggs.

Mike Levin's work

There is a lot of impressive work being done on chemical mechanisms controlling growth of organisms, including self-modifying or self-extending organisms of many sorts, including artificially created organisms. An important example is work done by Michael Levin's group on **xenobots**: <https://ase.tufts.edu/biology/labs/levin/publications/index.htm>

None of the examples I've seen so far approach the complexity, intricacy, and multiplicity of parallel developmental changes that occur during hatching of eggs of vertebrates. Moreover, as far as I can tell, he offers no explanation of how such chemical processes in eggs can produce complex cognitive skills, using knowledge that does not have to be (and could not be) derived from statistical evidence.

There is a huge amount of relevant literature, along with a growing number of online video presentations and tutorials, of which I have so far encountered only a small sample. So far I have not encountered any attempting to address the specific questions raised in this document about what chemical process in eggs of vertebrates can achieve, and how they do it.

Contrast with development in a uterus

Animals that develop within the mother's uterus, rather than inside an eggshell out of contact with internal parts of the mother, may be very different from animals that develop in the uterus, because the latter species allow much richer interactions between mother and foetus during the development of the foetus. In particular different chemical resources can be provided by the mother at different stages of development in the uterus, whereas the egg-laying mother has to provide all the chemical resources in the egg.

Information about the mammalian case is provided in a recent publication by Anna Ciaunica and colleagues:

The first prior: From co-embodiment to co-homeostasis in early life

Anna Ciaunica, Axel Constant, Hubert Preissl and Katerina Fotopoulou, in *Consciousness and Cognition* Volume 91, May 2021.

Note added 22 Sep 2022: Book by Jonathan Bard

I have recently been informed about this book *Evolution: The Origins and Mechanisms of Diversity* <https://www.amazon.co.uk/Evolution-Mechanisms-Diversity-Jonathan-Bard-ebook/dp/B09NQS91MX/> (published in 2021). From a partial reading, it appears to be very relevant to the above ideas, though I have not yet had time to decide whether there are any conflicts, or whether the book answers questions raised here. My impression so far is that it focuses on different questions. But a

close reading may reveal important connections, in which case details will be added here.

Some older work relevant to the Meta-Morphogenesis project

Jackie Chappell and Aaron Sloman, (2007), Natural and artificial meta-configured altricial information-processing systems, *International Journal of Unconventional Computing*, 3, 3, pp. 211--239,

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

Zsuzsanna Dancso University of Sydney Mathematics tutorials

The 2012 commemorative book mentioned above is:

Alan Turing: His Work and Impact, published in 2013,

Eds. S. Barry Cooper and J. van Leeuwen,

Publisher: Elsevier, Amsterdam,

It includes three invited papers presenting precursors of some of the ideas presented above.

The new ideas presented here are closely related to themes I have been exploring since I switched from research in mathematics and mathematical logic to research in philosophy of mathematics, around 1959.

Old publications that provide some of the ideas that are relevant to or motivate the most recent work include the following:

- Daniel C. Dennett, *Darwin's Dangerous Idea: Evolution and the Meanings of Life*, 1995, Penguin Press, London and New York,
- Lynn Margulis, *The Symbiotic Planet: A New Look at Evolution*, 1998, Weidenfeld and Nicholson, London
- Lynn Margulis and Dorion Sagan, *What is life?*, 1995, Simon and Schuster, New York,
- Erwin Schroedinger, *What is life?* 1944, CUP, Cambridge,
(Later reprinted in a different format, with additions.)
- Aaron Sloman, *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind*, Harvester Press (and Humanities Press), 1978, Hassocks, Sussex,
Now freely available online (html and PDF), with additional notes and some corrections:
<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp>
- A. Sloman, 'Necessary', 'A Priori' and 'Analytic', in *Analysis*, 1965, 26, 1, pp. 12--16,
<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1965-02>
(Based on my DPhil Thesis, 1962)
- Virtual-Machine Functionalism
(The only form of functionalism worth taking seriously in Philosophy of Mind and theories of Consciousness)
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>

- What About Their Internal Languages?

Commentary on three articles by Premack, D., Woodruff, G., by Griffin, D.R., and by Savage-Rumbaugh, E.S., Rumbaugh, D.R., Boysen, S. in *Behavioral and Brain Sciences Journal* 1978, 1 (4),

<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1978-02>

- Towards a grammar of emotions, in *New Universities Quarterly* 1982, 36, 3, pp. 230--238,

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

Hermann Rahn, Amos Ar and Charles V. Paganelli, (1979), How Bird Eggs Breathe, *Scientific American* 240, 2, Feb, 1979 pp. 46--55,

<https://www.jstor.org/stable/24965119>

NOTE

Comments, criticisms and suggestions are welcome, including comments pointing out that I am seriously mistaken! ([Email a dot sloman at bham.ac.uk](mailto:a.sloman@bham.ac.uk))

This work, and everything else on my website, unless explicitly excluded, is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).



If you use or comment on my ideas please include a URL if possible, so that readers can see the latest version.

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