DRAFT Oct 2019

BIOLOGICAL EVOLUTION'S USE OF REPRESENTATIONAL REDESCRIPTION

Aaron Sloman

School of Computer Science, University of Birmingham, UK

`http://www.cs.bham.ac.uk/~axs`

Introduction

Annette Karmiloff-Smith and I nearly became close colleagues when she was offered a chair at Sussex University. I wonder how different this tribute would have been, if we had interacted more closely over a longer period. I'll try to summarise some of the similarities and differences between our interests and our proposed explanatory theories. We were both deeply influenced by Jean Piaget, but in different ways. In particular some of Piaget's thinking, like much of my work was closely related to Immanuel Kant's ideas about the nature of mathematical discovery, an interest Annette did not share, though I'll try to show that some of her work on child development was relevant to mathematical cognition. We both drew attention to largely unnoticed complexities of gene-expression – but they were different complexities. There are implications for psychology, neuroscience, philosophy, AI, and theories of biological evolution.

After she moved to London, we met a few times at workshops and conferences, including an occasion when she invited me to London, to talk about representation. I think we may have discussed some of the ideas in (Karmiloff-Smith 1990). We also both contributed to a book on "Forms of Representation" (Peterson 1996), and in 2007, she was one of the guest speakers at a cognitive robotics project meeting in Paris in 2007, which I helped to organise.[1] But our interactions were rare and limited. We last met when I gave a departmental colloquium at Birkbeck in 2012 on "Toddler Theorems, Representational Redescription and Meta-Morphogenesis", and we talked after the seminar. She seemed to be disappointed that although I specified some of what needed to be explained, criticising over-simple alternative views, I was not able to demonstrate or report on a working AI system with the right explanatory powers. I think she had hoped that I would demonstrate or at least describe a neural-net based solution to the problem.

Despite the overlap in our ideas about cognitive development, described below, we had different centres of interest, and although I found her ideas and theories relevant to what I was trying to do, I do not think the interest was reciprocated, partly, I suspect, because I did not share her confidence in neural networks. We also had different but overlapping views of the

---

[1] `http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/latest.html`

relationship between the genome and development, a difference that I think I understand better as a result of writing this paper. In particular, I think Annette viewed epigenesis as a process of richly interacting branching processes, whereas the Meta-Configured genome theory, explained below, treats the genome as a multi-layered structure, in which some of the more recently evolved higher-level layers are expressed later, using information acquired from interactions between the earlier layers and the environment as parameters. This form of epigenesis, allows dramatic differences between results of gene expression in different environments, that could not result from any known type of uniform learning process. In particular, it allows a genome to specify discrete (parametrised) developmental stages, possibly related to Annette's examples of “representational redescription” mentioned below.

The rest of this chapter provides some (Kantian) background, then attempts to present some ideas about gene expression (the meta-configured genome) that seem to be capable of explaining Kant's observations, e.g. about mathematical discovery, which I'll relate briefly to Piaget's ideas, Annette's theories of representational redescription and more generally to current theories in psychology, neuroscience and AI/Computational cognitive science.

Background: Immanuel Kant on mathematical knowledge

A major theme in my research is trying to explain and extend what was correct in Kant’s philosophy of mathematics in (1781). Kant argued that mathematical knowledge, including knowledge of Euclidean geometry and elementary arithmetic, had features that distinguish it from both of the two kinds of knowledge identified by Hume, namely *analytic* knowledge (“relations of ideas”), based only on definitions of terms plus purely logical reasoning, and *empirical* knowledge (“matters of fact and real existence”) based on observation and experiment. Kant pointed out that besides being non-analytic (synthetic) and non-empirical (*apriori*), mathematical knowledge also had a *modal* feature: it was concerned with what is *possible* (e.g. polygons with N sides are possible for any N greater than 2), or *impossible* (e.g. two straight lines cannot completely enclose a finite portion of a plane) or *necessarily* the case: for example, angles of a planar triangle necessarily sum to half a rotation, i.e. 180 degrees[2]. Many such mathematical discoveries, some too trivial to be included in formal courses on mathematics, are unwittingly made and used in everyday intelligent actions by young children, without recognising their nature. I call them “toddler theorems”.[3]

My doctoral thesis (Sloman, 1962) defended Kant against 20th Century claims that his theories had been refuted by the success of Einstein’s General Theory of Relativity, which treated physical space as non-Euclidean. However, my defence of Kant did not propose *mechanisms* explaining ancient human mathematical abilities in geometry, topology and arithmetic. Later, after being introduced to Artificial Intelligence by Max Clowes, in 1969, and learning to program, I hoped to substantiate Kant’s claims by building working models of the relevant ancient geometrical discovery processes. Since mathematical competences of a six-year-old child are clearly not present at birth we need to find mechanisms that can produce mathematical knowledge of the kind Kant discussed, that is not available from birth and does not merely express generalisations from experience. I hoped to show how to design a “baby” robot that could learn about structures, processes and causation in something like the way a human child seems to, including going through processes of empirical discovery in various kinds of play and

---

[2] A summary of Kant’s claims about mathematics is in (Sloman 1965) and `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kant-maths.html`

[3] For examples, see `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html`

exploration, and then somehow coming to grasp deep non-empirical necessary truths, often without realising that that has happened.

Jean Piaget, Annette’s early mentor, had read Kant and Frege and by 1952 had made interesting discoveries about the development of mathematical cognition, including, for example, evidence suggesting that before the fifth year children could not grasp one of the facts that are crucial to understanding the natural numbers, namely that one-to-one correspondence is necessarily a transitive relation (a fact that is often distorted using the label “conservation”). However, Piaget was unable to propose mechanisms capable of explaining such mathematical insights.

He tried to find ways of modelling cognitive transitions using ideas from Boolean algebra and group theory, but they were not really relevant, and he apparently did not encounter computer programming and Artificial Intelligence until near the end of his life, when it was too late for him. If he had learnt to use a collection of important programming techniques half a century before they were developed in AI and software engineering, the history of developmental psychology (along with cognitive science and neuroscience) might have been very different. He commented on that in a talk he gave (in his wheelchair) at a conference in Geneva on Genetic Epistemology and AI in 1980, to which I had been invited.

Re-representing the precocial-altricial distinction

Informal observation of young children, squirrels and magpies in our garden, and reports and video recordings of other animals led me, during the 1990s, to conjecture that Kantian mathematical reasoning competences were related to the distinction between species with adequate competences available soon after birth or hatching (e.g. chicks able to peck for food and follow a hen, or deer able to walk to the mother’s nipple shortly after birth) and species that started off helpless but later achieved deeper and richer forms of intelligence, including humans, hunting mammals, squirrels, elephants and some birds. In the early 1990s I learnt that I was talking about the precocial/altricial distinction, long known to ethologists! Around that time, I also heard Manfred Spitzer give a talk emphasising the importance of delayed development of frontal lobes in humans – suggesting that evolution had discovered that development of some parts of brains in altricial species need to be delayed until they have access to enough information provided by older parts of brains interacting with the environment. That clearly had the benefit of allowing genetic mechanisms producing sophisticated competences to tailor their products to features of the environment, allowing such species to live and learn in environments with different kinds of complexity, as proposed in the meta-configured genome theory, below.

My thinking was disrupted when Jackie Chappell joined the school of Biosciences in Birmingham in 2004. She was one of the authors of a headline-grabbing paper (Weir, Chappell & Kacelnik 2002) reporting on Betty, a New Caledonian crow who repeatedly made hooks from straight pieces of wire, and used them to extract a bucket of food from a vertical glass tube. I later discovered, from the online videos at the Oxford ecology lab, that in ten trials Betty had made usable hooks in at least four significantly different ways, suggesting that something more was going on than merely acquiring food.

Jackie had earlier studied some of Betty's other forms of problem-solving behaviours, all involving spatial reasoning. Could Betty's desire to explore alternative solutions to an already solved problem be related to some of the phenomena Annette found in children's drawings? Certainly, in both cases fairly deep knowledge of alternative possible spatial arrangements and rearrangements of objects was used. But there are many unanswered questions. How do brains represent sets of possible alternatives? How do they identify some combinations as impossible (e.g. *A inside B, B inside C, and C inside A*)? Recognising and eliminating impossibilities from a search for solutions is an important aspect of intelligence – saving enormous amounts of wasted effort (unlike a wasp repeatedly trying to get through a pane of glass).

In response to my vague conjecture about altricial species, Jackie suggested that the basic distinction was not between *species*, or even individuals, but between *competences*. For example, humans are labelled an altricial species because newborn infants are helpless and incompetent. Nevertheless, they are born with very important competences, including the ability to obtain nourishment by sucking and swallowing, which requires complex coordination of a collection of muscles – a precocial competence. So we recommended replacing the biologists' distinction between altricial and precocial *species* with a distinction between altricial and precocial *competences* (Sloman & Chappell 2005, Sloman & Chappell 2007, Chappell & Sloman 2007*b*). Later we changed our terminology and labelled the competences "pre-configured" and "meta-configured".

Our claim was that some genetic specifications are *parametrised*, with information gaps that can be filled (given parameters) during gene expression. Obvious examples include genetically specified physiological structures, such as bones or muscles, that change in size, weight, strength, etc. during individual development, in coordination with changes in associated mechanisms, including neural control mechanisms. For example, larger bones will usually require larger, stronger muscles. We extended this kind of variability from genetically specified *physical* structures to genetically specified *competences*, including sucking, chewing, swallowing, grasping, and various forms of sensory information processing.

Preconfigured competences are genetically specified, and any changing parameters they need are either intrinsically generated or continuously derived from the changing developmental environment, including other parts of the developing organism. I think that description fits many of the cases of neural development described by Annette, e.g. in Karmiloff-Smith (2006), where multiple processes of development occur in parallel, with mutual influences. Later capabilities are results of such parallel, interacting developmental processes, in which information from the environment can play an important role, e.g. determining which language is spoken, in what accent, etc.

In contrast, *meta-configured* competences result from abstract, parametrised, *genetic* specifications that may be activated at different stages of development, using parameters that depend on information acquired during earlier gene expression. For example, if some gene expression produces behaviours whose results are recorded and stored, and a later phase of gene expression makes significant use of those recorded results, then later genes whose expression uses the previously acquired information are meta-configured, in our sense. Information acquired at an earlier stage is recorded for use as parameters for meta-configured competences developed later. For example, development of spoken words can use previously acquired phoneme production competences, partly derived from surrounding speakers. The differences between the two accounts of development can be understood by comparing the Meta-configured genome idea with the ideas expressed in Karmiloff-Smith (2006), which states: “...modules could be the result of ontogenesis over developmental time, not its starting point”.

In contrast, competences associated with meta-configured genes are essentially partly abstract *meta-competences* specified in the genome, though not expressed in the earliest phases of development. Instead, the abstractions are instantiated at appropriate later stages, using as parameters information previously gained through use of *pre-configured competences* developed earlier, or using results of earlier phases of meta-configured gene expression. Human language development is an example of such multi-layered meta-configured gene expression, including spectacular cases of later developments substantially re-organising the information gained in earlier phases of language development, including coping with exceptions to previously learnt syntactic regularities. All this allows genetically specified meta-competences, competence-*patterns*, to be instantiated in widely varying ways in different individuals, in different physical, social, linguistic, or cultural contexts, as illustrated by historical changes in children's playthings across centuries. A spectacular example is the recent development of new ontologies and extended concepts of causation required for use of wirelessly connected electronic devices, now taken for granted in many cultures.

The key idea is that meta-configured competence-patterns are instantiated using (not necessarily numeric) *parameters*, that are not specified by the genome but instead extracted from results of earlier interactions with the environment during development of simpler competences while the individual is younger and, in many cases, less physically mature. Simpler forms of parameter-based genomic flexibility (e.g. required for control of body parts changing in size, weight, strength etc.) may have been an evolutionary precursor of meta-configured genomes that take account of far more variability than that produced simply by individual growth and development in a standard environment.

Meta-configured genome expression

Figure 1: The Meta-configured Genome

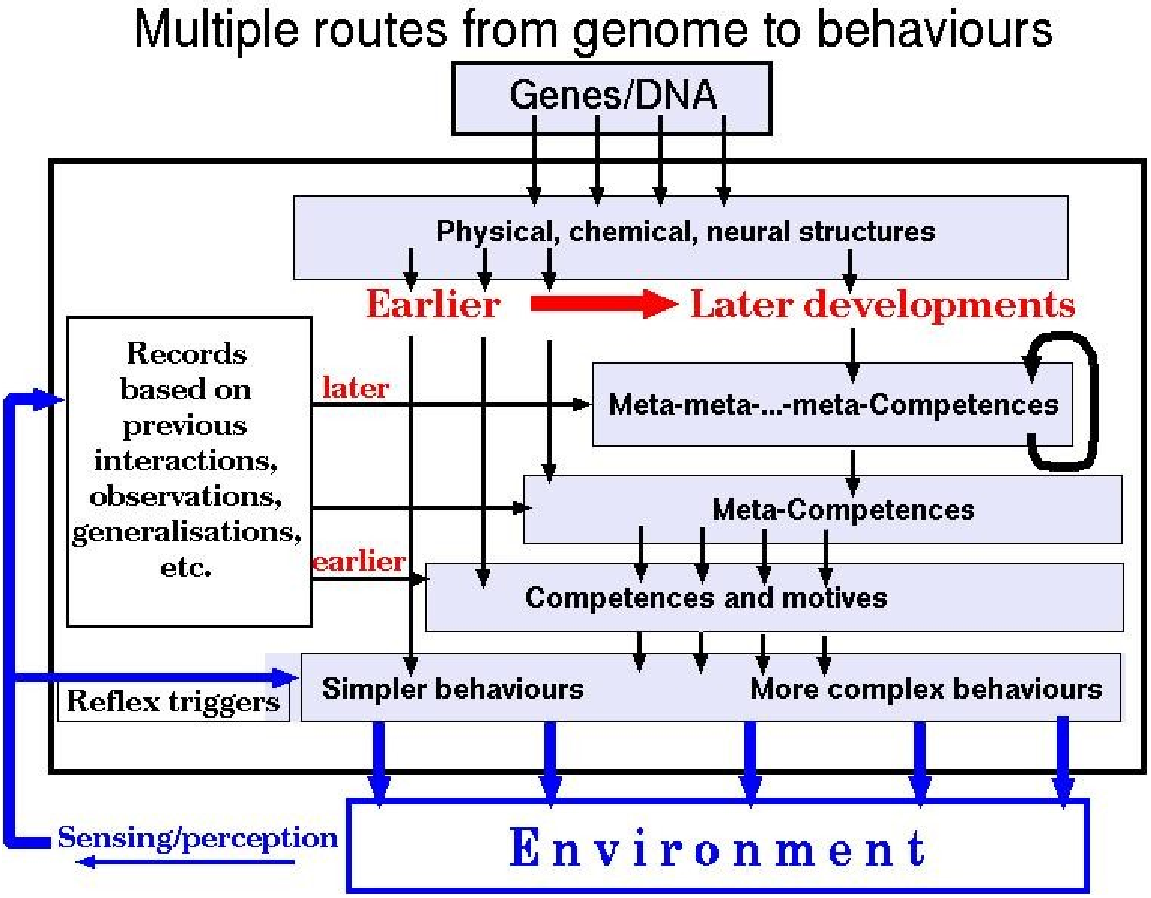

A meta-configured genome can produce meta-configured competences that are parametrised by information acquired using results of earlier gene expression influenced by the environment.

The diagram in Figure 1 crudely summarises processes associated with a meta-configured genome. (A video explaining how to read the diagram is available at https://www.youtube.com/watch?v=G8jNdBCAxVQ.) Downward arrows indicate multiple processes of gene-expression, starting with early, relatively direct, gene expression (down arrows to the left) and later gene expression (down arrows to the right), where later genes have parameters/gaps that are 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). 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.

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 this are likely to be highly sub-optimal. A teaching strategy that produces high scores on a particular test 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 interfering with development of deeper forms of text comprehension and creative thinking whose effects cannot be measured until much later.

Human language development provides spectacular examples of epigenetic variation based on a common genome. Several thousand different languages, including signed languages, have emerged, produced by communities sharing roughly the same language creation genetic mechanisms. The common human genome machinery enables production of thousands of different languages without specifying any of them! Similar comments apply to products of human technology at different times and locations. Even in some non-human species, the variety of structured percepts, intentions, intentional behaviours, information acquired (e.g. about locations of various resources and dangers, and techniques for achieving goals), would be impossible without the use of structured *internal* languages with compositional semantics, required, for example, for representing the unique, and changing, configuration of 3D obstacles, routes, potential perches, types of tool, control of body parts manipulating tools and the objects to which the tools are applied. Betty first used a newly made hook to lift the handle of a bucket.

A hive or colony of insects without this sort of meta-configured intelligence can instead make use of pheromone trails left by foraging individuals, but that is not available to animals that do not cooperate in such large numbers and which move around on complex extended terrain. Moreover, chemical trails may not survive strong winds and rain, and can only be used transiently by flying animals. So there are advantages in being able to reduce dependence on continuously available external cues and instead use enduring *internal* brain structures storing information about those surfaces. There are also enormous hidden costs in the complexity of mechanisms required for acquiring, transforming, combining, storing and using such information. That is one reason why brains are so complex.

Use of internal languages

All of this suggests that use of *internal* structured languages with compositional semantics evolved long before use of language for communication (Sloman 1978*b*). It is hard to explain some of the abilities of pre-verbal humans and other intelligent animals without postulating mechanisms for creating and manipulating internal information structures with compositional semantics, including structures that change both when perceived structures and relationships in the environment change, and also during consideration of *possible* actions or events and their consequences.

If such internal language use evolved first, then genetically specified mechanisms that originally made rich *internal* languages possible in a variety of species might later have been copied, modified and used to provide *external* languages. After some mechanism is copied for a new use, the new version can undergo changes that produce entirely new features. I suggest that originally communication between humans used a form of *sign* language, growing out of movements involved in cooperative actions (e.g. indicating direction of movement when cooperatively lifting and moving a heavy object), later extended to include *sounds* when hands were occupied or for communication to someone out of sight.

Such internal languages could support causal reasoning based on understanding of structural relationships rather than mere observed correlations – another product of a meta-configured genome – for example, understanding why a meshed pair of toothed wheels *must* rotate in opposite directions. As Piaget understood, this contradicts the widely shared (Humean) hypothesis that all causal reasoning is based on empirically learned generalisations, weighted by probabilities derived from recorded frequencies. Creative causal problem solving by crows, squirrels, pre-verbal humans, and other intelligent animals, renders statistical inference unnecessary for understanding some causal relationships.

After reading *Beyond Modularity* (Karmiloff-Smith 1992), I felt that its claims about “Representational Redescription” were closely related to our work on the meta-configured genome, and also some ideas about how results of exploration and experiment in a rich environment could feed into mathematical discoveries about the nature of the environment, illustrating Kant's claims in (1781) that the mathematical knowledge of necessary truths and impossibilities is (a) synthetic (not based on logical deductions from definitions) (b) *a priori* (non-empirical – not *derived from* nor *justified by* but *awakened by* experience), and (c) involves *necessary truths* and *impossibilities*, rather than degrees of probability.

Piaget on possibility and necessity

Annette did not share Piaget's interest in the Kantian questions that drove my research, including how humans could discover and make use of necessary connections and impossibilities, the topics of his last two (posthumously published) books (Piaget 1981, Piaget 1983). Annette had not read them, but it was she who informed me of their existence in 2010, after I asked her, by email, if she knew of any developmental or neural research on recognition of necessity or impossibility. She wrote back: *"Piaget devoted a whole year of the Centre d'Epistemologie Genetique to the child's becoming aware of necessity and not just empirical likelihood – 'it has to be so' – if you do a google there should be a symposium book on necessity"*. I followed her advice and acquired the two books arising out of that work, one on possibility and one on necessity. I found the experimental data fascinating but Piaget's explanatory theories too obscure and imprecise – probably because he had not learnt to produce computational explanations. Piaget and his collaborators in that research could not have been aware of Annette's work reported some time later, on children exploring possibilities by modifying drawings, and she seems to have been unaware of his work on possibility in which children were asked to specify new possibilities by modifying configurations of objects rather than by producing new drawings. I will not try to speculate about how each might have viewed the relationships.

My own work in AI was originally triggered by being challenged in 1971 to write a critical response to the defence of logic-based AI in (McCarthy & Hayes 1969). On that approach, all complexity of representation comes from the application of functions to arguments, as in logic and algebra, which I labelled use of "Fregean" representations, contrasted with "analogical representations" where properties and relations represent (not necessarily identical) properties and relations, as in 2D projections of 3D scenes (Sloman 1971). Analogical forms of representation play important roles in some of Annette’s examples of “representational redescription”, but her interest in those examples was different from mine. Perhaps our deepest disagreement was about mechanisms: I seek any AI techniques or mechanisms capable of explaining ancient mathematical discoveries and spatial reasoning in pre-verbal children, whereas she was particularly interested in general features of neural development, and how they contradicted some popular beliefs about how genes work, illustrated, for example in Williams syndrome. As a result, she seemed to me to be interested only in neurally inspired AI. I do not know whether she ever considered the possibility that much neural computation uses changing sub-neural chemical structures, with a mixture of continuous and discrete interactions rather than being restricted to changes in signal strengths and synaptic weights. Her emphasis on roles of brain chemistry in brain development might eventually have led her in that direction.

Overlaps

Annette and I shared some deep partly overlapping interests concerning what needs to be explained. In particular, we have both thought about *competence domains* but in different ways. I originally encountered this idea with the label “microworld” used by AI researchers, including (Minsky & Papert 1971) and I used it in teaching programming as well as characterising research problems (Sloman 1984). Later, I switched to using “domain” partly because of the overlap with ideas in *Beyond Modularity*, and partly because “micro-” suggests a restriction of scope. However, much of my thinking, unlike Annette’s, has been concerned with trying to understand Kantian aspects of development of mathematical competences: how can learners acquire and use knowledge of non-definitional necessary truths and impossibilities, rather than having to rely only on empirical generalisations where nothing is known unless based on observation and experiment.

Mathematical domains are implicit in some of Annette’s work, for example research on *balancing* tasks discussed in (Karmiloff-Smith & Inhelder 1974-1975). But as far as I know, like most developmental psychologists studying number cognition or spatial cognition, apart from Piaget, she never had the Kantian goal of trying to understand reasoning about necessary truths and impossibilities.

We were both interested in relationships between genetic mechanisms and development, but we disagreed about likely mechanisms insofar as she hoped that neural net models could explain all the phenomena (when we last talked in 2012), whereas I think that probabilistic mechanisms cannot explain mathematical understanding involving necessity and impossibility. Those are not extreme points on a probability scale. Despite such differences, I learnt a great deal from her work; I also found some of her recorded lectures on the complexity and diversity of developmental processes after fertilization inspiring, for example her 2011 talk at the National Academy of Sciences Sackler Colloquium on “Biological Embedding of Early Social Adversity: From Fruit Flies to Kindergartners”.

Representational Redescription

Annette’s work on “Representational Redescription” turned out to have interesting overlaps with some of my work on linguistic and mathematical development, including abilities to discover necessary connections and impossibilities, and the ideas concerning the meta-configured genome, mentioned above. However, as far as I know, Annette did not explicitly emphasise, as Jackie Chappell and I did, that new types of learning competence can arise out of the interplay between the “meta-configured” genetic mechanisms expressed relatively late in development, and what has previously been learnt. For example, a nine-year-old is capable of learning things a toddler cannot, and some of that may be dependent on having learnt ways of learning that are relevant to the current environment. Another example would be 21st Century five-year-olds learning how to use a computational system with a graphical interface, which none of their ancestors ever had a chance to do.

Such examples illustrate the fact that cultural evolution can feed information into the products of biological evolution in ways that effect the learning and developmental competences produced by the genome, but without altering the genome. This depends on the later-developing learning mechanisms produced by evolution not being fully specified. In effect, they are parametrised and the parameters can be filled by products of earlier learning. This seems to be consistent with the general thesis of *Beyond Modularity*. The discussion of language learning in Beyond Modularity, which involves differences between cultures and between learning a spoken language and a sign language, implicitly makes the point. Such learning depends on the parameter-based meta-configured genome, which allows parameters acquired at an earlier stage of development to feed into a template “turned on” by the genome at a later stage, as illustrated by the Meta-Configured genome diagram above. For example, a template relevant to grammatical structures may obtain inputs acquired from spoken, written, or sign languages.

I do not know whether any of Annette’s ideas changed after we last met or whether she would have been interested in my attempts to fill gaps in current explanations of ancient mathematical competences. I know of nobody with good explanations of how ancient mathematical brains made possible the amazing discoveries of Archimedes, Euclid, Zeno and many others. I suspect they used mechanisms that had previously enabled more general discoveries about spatial impossibilities and necessities made much earlier, using mechanisms shared with many other species with deep spatial intelligence, though without human reflective (meta-cognitive) abilities.

Neither current symbolic AI, nor artificial neural nets, nor known brain mechanisms seem to have the capabilities required to explain mathematical discoveries of the sorts described by Kant, including representing “alethic” modalities, such as *impossibility* and *necessity*. These have nothing to do with probabilities or statistical evidence. Moreover, the ancient examples are not cases of logical/definitional impossibility or necessity, but involve discovery of necessary features of spatial structures, such as transitivity of containment, or the impossibility of intersection of two convex 2D or 3D shapes producing a non-convex shape. At the time of writing I do not know of any good account of the workings of ancient mathematical minds, or even the spatial reasoning in squirrels, weaver birds, elephants, octopuses, apes, pre-verbal toddlers, and other animals that are good (but not perfect) at spatial reasoning. Neither are there working robots with those abilities. I do not think these abilities can be innate (neither did Kant). Rather their discovery involves something like what Annette called representational redescription, and Kant called synthetic necessary truths discovered non-empirically: not derived from experience but “awakened” by experience.

Much work in artificial intelligence attempts to design machines that start with very little world knowledge but include a powerful learning mechanism that enables such machines, possibly helped by teachers who set challenges and indicate whether responses to questions and challenges are correct or not. The hope of such research is that over time the learning mechanism will discover statistical relationships between features in its sensory and motor records that allow discoveries to be made concerning how to predict what is likely to follow some collection of actions and sensed features, and on that basis will reliably allow intelligent selection of actions to achieve desired results. Like Kant, I do not believe that such statistical competences can explain mathematical discoveries in geometry and topology – because noticing statistical regularities cannot lead to the ancient mathematical discoveries about *necessary* relationships. Those discoveries require special mechanisms. Perhaps they grew out of mechanisms that evolved because they were powerful tools for spatial reasoning, like the use of diagrams (rather than logical formulae) in constructing geometric proofs, since ancient times?

Numbers

I think much of what psychologists and neuroscientists have written about development of number competences is wrong, or at least seriously incomplete, because they are based on an incorrect analysis of what numbers (of various kinds) are. That is not surprising – philosophy of mathematics is a complex subject and is not normally taught in those disciplines.

One use of numbers by scientists and engineers is in measurements of length, area, volume, weight, force and a host of other properties of physical objects, such as tension, torsion, momentum, and energy. Some animals have behaviours that seem to include very precise numerical measures, for example spider monkeys swinging on, bouncing on, and leaping between structures, including performing several such actions in quick succession, implying both very rapid and accurate measurements and very rapid and accurate control of muscular tension and compression, producing forces that accurately launch them through space onto a rigid or flexible target object. Presumably a combination of an extended period of evolution, combined with learning processes during development produces a tightly integrated collection of “compiled” skills, but with very little reflective understanding of what those skills are and how they interact.

Such competences, including all the acquired measures, may be available for immediate use during performance of actions (e.g. in spider monkeys), without ever becoming objects of reflective thought, for instance thoughts about relationships between the elasticity of a launching support, the distance to a target branch, the muscular forces needed to produce an accurate launch, etc. In contrast, a human with vastly inferior versions of those skills (like me) can notice them, think about them, and ask questions about them without being able to acquire and use them (trapeze artists are rare exceptions.) That suggests that the brain mechanisms required for the different uses of “numerical” measures are very different and may take different forms in different evolutionary lineages. I suspect similar comments are applicable to tests for “numeracy” in various non-human animals and in young children.

This line of thought makes me deeply suspicious of research reports claiming to have identified numerical/mathematical competences in infants or toddlers and other animals when nobody knows exactly what sorts of brain functions are actually being used in the various cases. I suspect similar comments can be made regarding the discussion of “The child as a mathematician” in Beyond Modularity. If our natural number concepts correspond to equivalence classes based on 1-1 correspondence, as agreed by David Hume, Gottlob Frege and Bertrand Russell, then many of the behavioural tests alleged to identify onset of number competences may be irrelevant if they merely produce evidence for competencies that produce the right results in a very limited set of contexts, just as the jumping precision of a spider monkey happens to correspond to use of the physicists concepts of length, force, mass, and momentum, without justifying the assumption that anything in the monkey’s brain uses the same concepts.

Example mechanisms capable of detecting or creating one-to-one correspondences in various practical tasks were presented in (Sloman 1978*a*), which assumed the existence of our current number and counting systems (based on one-to-one correspondence), with tentatively sketched mechanisms capable of supporting their use (I here refer not to physical (e.g. brain) mechanisms, but to computational (virtual machine) mechanisms that could be implemented in different physical mechanisms, including mechanisms capable of generating parallel, coordinated, discrete processes, such as pointing while counting). I do not think there is any empirical or neural evidence that human brains are innately programmed with abilities to think about and reason about 1-1 correspondences involving arbitrarily large sets. The relevant abilities must have developed through cultural processes, building on more general innate or learnt capabilities. For understanding numbers (or geometry), however, each individual has to develop insight into impossibilities and necessary connections.

Beyond Modularity also emphasises the fact that much learning that requires innate mechanisms can have features that are strongly dependent on what it is about the environment that is learnt; for example, learning about properties of physical structures and processes that vary widely across cultures – such as building materials and what they are used for. Such learning can occur at different stages of development, and still be heavily influenced or constrained by the genome, in addition to being influenced by portions of the environment acted on and perceived – for example, abilities to write, understand, test and debug computer programs, for which opportunities did not exist until very recently. It is likely that among known animals only humans have genes that have the ability to support learning to debug computer programs, and have had them for millennia. Some other species may have related but simpler abilities to debug action strategies.

Chapter 4 of Beyond Modularity "The Child as a Mathematician" requires critical discussion: firstly, because learning about numbers is a much more complex and diverse process than the chapter acknowledges; and secondly, because there is much more to mathematical development, including toddler mathematical development, than learning about numbers. For example, as Piaget understood, it should include learning about topological relationships, including containment and connectivity.

Nevertheless, *Beyond Modularity* (or an updated version) should be compulsory reading for all researchers working in the areas of intelligent robotics, cognitive robotics, AI learning systems, and more generally computational cognitive science. Reading the book will give such researchers important new insights into what needs to be modelled and explained as well as experience of debugging a complex theory!

I do not agree with everything in the book. In particular, I believe that the emphasis only on *representational* change, without discussing required changes in information processing *architectures and mechanisms* during development, is a major gap. In her opposition to modularity theories that postulate specific innate competences, Annette claims that the processes that happen within a domain to produce observed developments use general mechanisms. This is not correct, because learning to learn, as discussed in (Chappell & Sloman 2007*b*), can include learning to learn different kinds of skills, facts, and learning strategies relevant to different environments or different developmental stages. Towards the end of Beyond Modularity Annette raises doubts about the adequacy of the concepts and theories presented. It is possible that later work stimulated by her ideas has made important progress about which I am still ignorant.

Evolution’s Use of Representational Redescription

I have been working on evolution’s use of compositionality and its role in ancient mathematical discovery, which I think is at least loosely related to Annette’s ideas. Roughly: evolution discovered increasingly powerful abstractions that can contribute to increasingly powerful modes of compositionality in its designs, by means that cannot be replicated in individual learning and development. Some of the results were the ancient mathematical discoveries that are now in constant use by scientists, engineers, architects etc., but which are not explicable by current neuroscience, nor replicated in AI systems. Those evolutionary discoveries, whose precise history remains to be charted, have some of the character of processes of representational redescription Annette attributed to individuals. But I am attributing them to evolutionary processes. I cannot tell whether she would have approved.

By encoding abstracted versions of new discoveries in genomes, evolution left mechanisms of development in individuals the task of instantiating and combining the abstractions, not re-discovering them, though differences in individual trajectories resulting from layered, context-sensitive, stages of instantiation, may produce amazing individual creativity. If what is in the genome is appropriately abstract, then it can go on being applied to new products of cultural evolution instead of being restricted to the original applications. Linguistic examples would include late-expressed genes that deal with differences between singular and plural forms, or differences in tenses. If the information is not encoded in terms of the precise language in use when the new capability is added to the genome then it can be useful in a far greater variety of future environments. This can apply to semantic and pragmatic features of language, and to internal uses of language for reasoning, as well as to grammatical details.

There are many striking examples in the (presumably) common collection of parametrised genomic features that can be instantiated in different ways to produce thousands of different languages, differing in primitive sounds or signs used, in syntactic forms, in semantic contents, in pragmatic communicative functions, and in various social processes. For example, the ability to congratulate someone on achieving something admirable but difficult is unlikely to have direct genetic support yet can be expressed in multiple different languages by combining more primitive linguistic functions, a possibility that can, in principle be re-discovered in different cultures, or by different individuals. Moreover, that possibility must have been supported by the human genome long before any human ever made use of it.

Among the products of all those evolutionary mechanisms are results that cannot be produced by current models of deep learning or known neural mechanisms, nor by purely logical/algebraic reasoning, for example the achievements of ancient geometers mentioned above. A major task is to specify *detailed requirements* for the mechanisms underlying ancient mathematical discoveries, and other aspects of natural intelligence, a task that I do not think has ever been done systematically, comprehensively and adequately, although *Beyond Modularity* contributes important relevant ideas. In particular, we need a theory about the internal languages required, not for communication with other individuals, but for encoding/representing information contents for *internal* use, including: perceptual contents, emotional states, online action control, offline consideration of actual or possible actions, and reflection on these internal representations.

Conclusion

Originally inspired by Kant's philosophy of mathematics, I have been working for some time on evolution's use of compositionality and its role in ancient mathematical discovery, which I think is related to Annette's ideas. Roughly: evolution discovered increasingly powerful abstractions that can contribute to increasingly powerful modes of compositionality by means that cannot be replicated in individual learning. So those abilities need to be supported genetically. In other words, powerful cognitive abilities depend on genetically determined modules, but some of the most important modules are not active from birth, and they are only partially specified abstractions, with important parameters acquired from the environment before the modules are activated.

By encoding abstracted versions of those discoveries in genomes, evolution provided opportunities for individuals to solve difficult problems, including engineering design problems, mathematical problems and problems of explaining new scientific observations, by instantiating powerful abstractions, without having to re-discover them, though with differences in individual trajectories resulting from layered, context-sensitive, results of such instantiation. This can produce amazing individual creativity, illustrated, for example, in Annette's work on children's drawings. Results based on understanding spatial necessities/impossibilities, cannot be produced by current models of deep learning or known neural mechanisms, since those mechanisms cannot represent, let alone discover, impossibilities and necessities.. Neither do powerful current logic-based AI reasoning systems capture the ancient mathematical discovery processes using spatial reasoning, or the closely related spatial reasoning capabilities of pre-verbal toddlers and other intelligent animals.

I felt very honoured, and very surprised, when I was invited to contribute to this book. However, reading or re-reading both things Annette wrote and work by others reporting, commenting on or criticising her work, has helped me to understand better the depth and breadth of her work, and some of the details of our agreement and disagreement. I hope some of the relationships summarised briefly here will trigger interest in new research that addresses important scientific and educational problems, as it has prompted me to clarify some aspects of the Meta-Configured Genome project. Thank you Annette!

References

Chappell, J. & Sloman, A. 2007*a*, ‘Contributions to WONAC: International Workshop on Natural and Artificial Cognition Two ways of understanding causation: Humean and Kantian’.

Chappell, J. & Sloman, A. 2007*b*, ‘Natural and artificial meta-configured altricial information-processing systems’, *International Journal of Unconventional Computing* **3**(3), 211–239.

Clowes, M. 1971, ‘On seeing things’, *Artificial Intelligence* **2**(1), 79–116.

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

Fuson, K. C. 1979, ‘Review of *The Child’s Understanding of Number by R. Gelman and C. R. Gallistel*’, *Journal for Research in Mathematics Education* **10,**(5), 383–387.

Gelman, R. & Gallistel, C. R. 1978, *The Child’s Understanding of Number*, Harvard University Press.

Hilbert, D. 1899, *The Foundations of Geometry*, Project Gutenberg, Salt Lake City. Translated 1902 by E.J. Townsend, from 1899 German edition.

Huffman, D. A. 1971, Impossible objects as nonsense sentences, *in* B. Meltzer & D. Michie, eds, ‘Machine Intelligence’, Vol. 6, Edinburgh University Press, 22 George Square, Edinburgh, Great Britain, pp. 295–323.

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

Kant, I. 1781, *Critique of Pure Reason*, Macmillan, London. Translated (1929) by Norman Kemp Smith.

Karmiloff-Smith, A. 1990, ‘Constraints on representational change: Evidence from children’s drawing’, *Cognition* **34**, 57–83.

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

Karmiloff-Smith, A. & Inhelder, B. 1974-1975, ‘If you want to get ahead, get a theory’, *Cognition* **3**, 195–212. 3.

Karmiloff-Smith, A. 2006, The tortuous route from genes to behavior: A neuroconstructivist approach, Cognitive, affective & behavioral neuroscience, 6, 9--17

McCarthy, J. 2008, ‘The well-designed child’, *Artificial Intelligence* **172**(18), 2003–2014.

McCarthy, J. & Hayes, P. 1969, Some philosophical problems from the standpoint of AI, *in* B. Meltzer & D. Michie, eds, ‘Machine Intelligence 4’, Edinburgh University Press, Edinburgh, Scotland, pp. 463–502.

Minsky, M. & Papert, S. 1971, *Progress Report on Artificial Intelligence*, MIT. AI Memo AIM-252.

Mumford, D. 2016, ‘Grammar is not merely part of language’, http://www.dam.brown.edu/people/mumford/blog/2016/grammar.html. Online Blog.

Neisser, U. 1976, *Cognition and Reality*, W.H. Freeman, San Francisco.

Peters, L., Messer, D., Davey, N. & Smith, P. 1999, ‘An investigation into Karmiloff-Smith’s RR model: The effects of structured tuition’, *Brit Journ of Developmental Psychology* **17**, 277–292.

Peterson, D., ed. 1996, *Forms of representation: an interdisciplinary theme for cognitive science*, Intellect Books, Exeter, U.K.

Piaget, J. 1981, *Possibility and Necessity Vol 1. The role of possibility in cognitive development*, U. of Minnesota Press, Minneapolis. Tr. by Helga Feider from French in 1987.

Piaget, J. 1983, *Possibility and Necessity Vol 2. The role of necessity in cognitive development*, U. of Minnesota Press, Minneapolis. Tr. by Helga Feider from French in 1987.

Roberts, L. G. 1965, Machine perception of threedimensional solids, *in* J. T. et al., ed., ‘Optical and Electrooptical Infor mation Processing’, MIT Press, Cambridge, MA.

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

Sloman, A. 1962, Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil Thesis), PhD thesis, Oxford University.

Sloman, A. 1965, ‘‘Necessary’, ‘A Priori’ and ‘Analytic’’, *Analysis* **26**(1), 12–16.

Sloman, A. 1971, Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence, *in* ‘Proc 2nd IJCAI’, William Kaufmann, London, pp. 209–226. Reprinted in *Artificial Intelligence*, vol 2, 3-4, pp 209-225, 1971.

Sloman, A. 1978*a*, *The Computer Revolution in Philosophy*, Harvester Press (and Humanities Press), Hassocks, Sussex. http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp, Revised 2018.

Sloman, A. 1978*b*, ‘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)’, *Behavioral and Brain Sciences* **1**(4), 515.

Sloman, A. 1982, Computational Epistemology, *in* ‘Cahiers De La Fondation Archives Jean Piaget, No 2-3 Geneva June 1982’, Archives Jean Piaget. (Proceedings Seminar on Genetic Epistemology and Artificial Intelligence, Geneva 1980), http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#53.

Sloman, A. 1984, Beginners need powerful systems, *in* M. Yazdani, ed., ‘New horizons in educational computing’, Ellis Horwood Series In Artificial Intelligence, Chichester, pp. 220–234. http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#45.

Sloman, A. 2016, Natural Vision and Mathematics: Seeing Impossibilities, *in* U. Furbach & C. Schon, eds, ‘Second Workshop on: Bridging the Gap between Human and Automated Reasoning, at IJCAI 2016’, New York, pp. 86–101. http://ceur-ws.org/Vol-1651/.

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

Sloman, A. & Chappell, J. 2007, ‘Computational Cognitive Epigenetics (Commentary on (Jablonka & Lamb 2005))’, *Behavioral and Brain Sciences* **30**(4), 375–6.

Smith, L. 1994, ‘The Development of Modal Understanding: Piaget’s *Possibility* and *Necessity*’, *New Ideas in Psychology* **12,**(1,), 73–87.

Trettenbrein, P. C. 2016, ‘The Demise of the Synapse As the Locus of Memory: A Looming Paradigm Shift?’, *Frontiers in Systems Neuroscience*.

Turing, A. M. 1938, ‘Systems of Logic Based on Ordinals’, *Proc. London Mathematical Society* pp. 161–228.

Turing, A. M. 1950, ‘Computing machinery and intelligence’, *Mind* **59**, 433–460.

Turing, A. M. 1952, ‘The Chemical Basis Of Morphogenesis’, *Phil. Trans. R. Soc. London B 237* **237**, 37–72.

Weir, A. A. S., Chappell, J. & Kacelnik, A. 2002, ‘Shaping of hooks in New Caledonian crows’, *Science* 297(9 August 2002), 981.