Gaps Between Human and Artificial Mathematics

Aaron Sloman (a.sloman@cs.bham.ac.uk)
School of Computer Science, University of Birmingham,
Birmingham B15 2TT, UK

Abstract

The Turing-inspired Meta-morphogenesis project begun in 2011 was partly motivated by deep gaps in our understanding of mathematical cognition and other aspects of human and non-human intelligence and our inability to model them. The project attempts to identify previously unnoticed evolutionary transitions in biological information processing related to gaps in our current understanding of cognition. Analysis of such transitions may also shed light on gaps in current AI. This is very different from attempts to study human mathematical cognition directly, e.g. via observation, experiment, neural imaging, etc. Fashionable ideas about "embodied cognition", "enactivism", and "situated cognition", focus on shallow products of evolution, ignoring pressures to evolve increasingly *dis*embodied forms of cognition to meet increasingly complex and varied challenges produced by articulated physical forms, multiple sensory capabilities, geographical and temporal spread of important information and other resources, and "other-related meta-cognition" concerning mental states, processes and capabilities of other individuals. Computers are normally thought of as good at mathematics: they perform logical, arithmetical and statistical calculations and manipulate formulas, at enormous speeds, but still lack abilities in humans and other animals to perceive and understand geometrical and topological possibilities and constraints that (a) are required for perception and use of affordances, and (b) play roles in mathematical, and proto-mathematical, discoveries made by ancient mathematicians, human toddlers and other intelligent animals. Neurally inspired, statistics-based (e.g."deep learning") models cannot explain recognition and understanding of mathematical *necessity* or *impossibility*. A partial (neo-Kantian) analysis of types of evolved biological information processing capability still missing from our models may inspire new kinds of research helping to fill the gaps. Had Turing lived long enough to develop his ideas on morphogenesis, he might have done this.

Keywords: Archimedes; Euclid; Kant; geometry; topology; vision; evolution; biological information processing; limitations of current computational models evolution as a blind mathematician.

Introduction

There are deep gaps in current AI models, related to gaps in theories of cognition, especially mathematical cognition (despite impressive mathematical powers of computers). The Turing-inspired Meta-Morphogenesis project, proposed in 2011 asks new questions about evolution of biological information processing, identifying what needs to be explained and possible types of explanation corresponding to different evolutionary stages.[1] Large sums are being spent in the hope that more training on more data can diminish, and eventually remove, those gaps, guided by research on how humans acquire the relevant competences and on brain mechanisms involved, but the research focuses on a *subset* of the relevant competences and mechanisms, leaving much unexplained. E.g. research that focuses on numerical competences, ignores geometric and topological competences, that are arguably more fundamental, in ways that I'll explain later. Moreover research on statistics based learning cannot explain discoveries of necessary truths, e.g. geometrical, topological and arithmetic truths.

Many psychologists also ignore important mathematical features of competences being investigated, because they don't clearly distinguish empirical from non-empirical learning. For example, not all psychologists studying number cognition seem to realise that full understanding of cardinal and ordinal numbers depends on grasping that one-to-one correspondence (bijection) is a transitive and symmetric relation (and therefore also reflexive), and moreover those properties are *necessary* (i.e. non-contingent) features of bijection, but not logical or definitional features. This was pointed out by Kant in 1781, though he knew of no explanatory mechanisms. My 1962 thesis (now online) defended Kant against common criticisms, but I had never heard of AI then and I lacked the opportunity to base a defence on computational modelling, a gap I began trying to fill in my 1978 book. Four decades later there still seem to be no working AI systems able to replicate the discoveries in topology, geometry and arithmetic, made by ancient mathematicians such as Archimedes, Euclid, Zeno and others, nor the closely related, hard to observe, discoveries unwittingly made by pre-verbal human toddlers,[2] or even squirrels and nest-building birds.

A rich sample of approaches to the problems of characterising and explaining numerical competences can be found in a BBS survey by Rips *et.al.*, including commentaries and responses. Unfortunately influences on individual mathematical development now are so diverse, including biological, physical, cultural, educational and individual differences, and so little attention is paid to the problem of specifying *implementable mechanisms*, as opposed to verbal descriptions of what brains or minds do, that the research is inevitably fragmentary and inconclusive and proposed theories lack the precision required to guide designs for testable working models.

Piaget drew attention to many combinations of competence and incompetence displayed by children, and produced evidence that most did not understand that 1-1 correspondence is a transitive relation until they are five or six years old. It is also symmetric, unlike many transitive relations children learn about (e.g. "taller than", "heavier than"). Unfortunately, calling this learning about "conservation" misleadingly sug-

[1] References have been deleted in this version for lack of space, but can be found, with links to online papers, at http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-iccm17.pdf

[2] Like the 17.5 month old child apparently testing a conjecture in 3D topology here http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#pencil

gests that understanding preservation of numerosity across spatial rearrangement is related to understanding that matter is conserved when rearranged. One-to-one correspondences can hold between completely abstract entities that have no matter to conserve. This is obvious to mathematicians, but perhaps not all developmental researchers.

One common Piagetian test for understanding numerosity tends to use examples of two types (e.g. apples, bananas) and supertype (e.g. fruit) in situations where there are (e.g.) more apples than bananas and children are asked whether there are more apples or more fruit. At a certain age they tend to say "more apples". However, there is usually no attempt to check that they have understood the question as intended. I found that if a child is asked to count the apples, then to count the fruit, then asked the same question, the correct answer is given. Some then generalise, without help, to other cases, e.g. giving the right answer to the question "Are there more open windows or more windows?" asked of a building with far more windows open than shut. This suggests that some children interpret the original question wrongly. I don't know if any psychologist has tried tampering with Piaget's experiment in this way. However, Margaret Donaldson showed in 1978 that slight variants of some of Piaget's other experiments, produced significantly different results.

My aim is not to criticise Piaget or his (often less well informed) followers but to draw attention to problems of empirical research not based on deep theories. Is there a deep theory in neuroscience capable of explaining what sort of late developing neural mechanism can change the powers of a child's brain so that the necessary transitivity and symmetry of one-to-one correspondence is grasped? This can be viewed as a topological problem about two networks of connections, e.g. a network formed by setting up a one-to-one correspondence between elements of sets A and B, and one between elements of B and C. We can see (How?) that if A, B and C are disjoint sets, the two sets of links can always be concatenated to form one-to-one relationships between A and C. Does anyone have a theory as to how brain mechanisms can detect, or even represent, the impossibility of any counter-example – i.e. the fact that the transitivity is a *necessary* truth? The work of mathematical logicians (e.g. Frege, Russell and others) allows the transitivity to be proved (tediously) in a formal logical system, but it was understood by ancient mathematicians (and young learners), centuries before those formal proof methods had been discovered. What happened in their brains when the *necessity* of transitivity of bijection, i.e. the *impossibility* of counter examples, was grasped?

Mathematical discoveries are not concerned with empirical or contingent regularities but with necessary connections and impossibilities (e.g. internal angles of a planar triangle *necessarily* sum to half a rotation, and it is *impossible* for any number to be the largest prime). How could we check that a brain mechanism is able to represent and use these notions of necessity and impossibility, which are features of mathematical discoveries, but not empirical discoveries? The answer will depend in part on a good theory of the semantics of modal concepts – often taken nowadays to be "possible world" semantics.[3] However, ancient mathematicians did not need this notion of a possible world: they were exploring compatible and incompatible collections of relationships in *this* world, often represented diagrammatically (Sloman 1962). .

So mathematical (as opposed to empirical) discoveries about numbers, lines, angles, etc. require use of (alethic) *modal* concepts (e.g. "possible", "impossible", "necessarily true", "necessarily false"). Standard ways of acquiring general information by observing instances and collecting statistics, cannot yield such mathematical knowledge, since that requires more than observed regularities. Perhaps many badly taught learners never get beyond memorising what they have been taught, but that's not what needs explaining.

I am not aware of any computational model that is able to replicate not only those arithmetical and geometrical discoveries but also other topological impossibilities that children seem to understand without mathematical training, for example that two solid rings cannot become linked and unlinked simply by being moved continuously, or that a shoe-lace cannot be pulled out of lace-holes twice as fast by pulling both ends at once. Nor does any AI model that I know of explain this. There is no evidence that AI theorem provers that draw conclusions from logical axioms can model what a young child, or an intelligent squirrel or crow does, or what ancient mathematicians did over 23 centuries ago, long before discovery of modern logic and algebra, and Descartes' use of arithmetic to model geometry.

Kant pointed out that ancient mathematical discoveries are characterised by being (a) non-empirical, (b) non-analytic (i.e. not derivable from definitions using only logic) and (c) non-contingent – the truths and falsehoods are instances of *necessity* and *impossibility* as explained in my thesis. This does not imply that mathematicians are *infallible*: they can and do make mistakes of various sorts, though they often discover and correct their mistakes, as demonstrated in *Proofs and Refutations* by Lakatos.

The 20th Century discovery that physical space is non-Euclidean is often regarded as demonstrating that Kant was wrong about mathematical knowledge, whereas it merely shows that some of his examples were wrong. He could have used the discovery that a subset of Euclidean geometry can be extended in different ways, yielding Euclidean and non-Euclidean geometries, as an example of a mathematical truth that is synthetic, necessarily true and not empirically based. Non-Euclidean geometries had been discovered before the 1919 eclipse showed that physical space was not Euclidean. Such discoveries add to what needs to be explained by neuroscience and modelled by AI.

Regarding arithmetic: is there a neural theory explaining how brains generate and control parallel sequences of actions required in counting operations of various sorts, with different stopping conditions depending on the task and various

[3] https://en.wikipedia.org/wiki/Possible_world

ways in which counting errors can be detected and be corrected, as described in Ch. 8 of *CRP*?[4] An explanatory mechanism should explain how counting can be applied, via different senses and movable body parts, to events, continuous processes (e.g. rotations, changes of direction, skin strokes, or sound oscillations), to static objects, and to abstract entities (e.g. numbers, words), along with self-monitoring to detect departures from strict one-to-one correspondence. Moreover, some mathematical discoveries can be made by noticing novel features of such thinking processes, e.g. repeated patterns. I suspect no known *neural* mechanism explains how reflection on processes produced by number generating mechanisms can lead to the concept of a non-terminating sequence, and then to an understanding that there are infinitely many numbers. What allows a child to understand "never stops"?

Another discovery that I believe is beyond current AI theorem provers was known to Archimedes and others: namely adding the *neusis* construction to Euclidean geometry, allowing motion of a straight-edge with two marks, makes it easy to trisect an arbitrary angle, which is impossible in standard Euclidean geometry.[5] What would a neural explanation of such a discovery process look like? Finding brain regions that are active during such discoveries does not tell us how brains encode universally quantified semantic content, or how they derive new semantic contents. It cannot be assumed that such discoveries are based on applying rules of modern logic (e.g. predicate calculus) to logical axioms, in part because modern logic was not available to ancient mathematics: it was mostly created recently by thinkers like Boole, Peano, Frege, Russell and others. Moreover, Euclidean geometry was not axiomatised using modern logic until 1899, by David Hilbert. Trisection was proved impossible in that system, so discovery of a construction that trisects an arbitrary angle must have used a different mode of spatial reasoning. I suspect ancient discoveries in geometry and topology were closely related to the need to identify positive and negative affordances, shared with other intelligent species. But evolution added some additional, unknown(?) discovery or reasoning mechanism in humans.

Meta-cognitive mechanisms, allowing internal processes based on previous competences to become objects of reflection during their performance seem to be required for some new mathematical insights. Many practical tasks can make use of multiplication and division, e.g. making sure that every member of a group has two shoes, or dividing N tasks between M people. Reflecting on this leads to the discovery that some sets with N members can be divided into M equal sets, but not into M+1 equal sets, and eventually that some numbers *cannot* be divided into any number of equal sets: they are primes, already familiar to Euclid. It is not clear how the *impossibility* is recognized, as opposed to mere *repeated failure*. Statistics-based learning mechanisms could not discover impossibilities and necessary truths: those are not degrees of probability. (However mathematical theorems about probabilities are necessary truths, not probabilistic assertions!)

Piaget (who had studied Kant, Frege and Russell) understood some of the problems. His two posthumous books were on *possibility* and *necessity*, though he lacked the tools required to solve our problems.

Mathematical meta-cognition

Metacognitive reasoning processes seem to have enabled Euclid (or a predecessor) to discover and prove that there cannot be a largest prime number, so there must be infinitely many prime numbers. How did evolution produce mechanisms with such capabilities, and how do they work? Perhaps a "duplicate then differentiate" transition in our evolutionary history somehow produced meta-cognitive capabilities, allowing comparisons of modes of thinking on different occasions, leading to important insights concerning differences between reliable and unreliable reasoning, enabling introspected reasoning processes to be described and modified while they were being performed, and allowing mistakes of reasoning to be discovered and eliminated, or successful modes to be combined to form more complex modes.

Such meta-cognitive abilities would also have social consequences, e.g. allowing strategies discovered during self-debugging to be later taught to others.[6] Every good mathematics teacher knows that learning to detect mistakes in reasoning is a deep part of mathematical education. More generally, the extension of meta-cognition from direct self-observation to indirect other-observation can help with effective other-debugging processes. I don't know if anyone has an appropriately deep theory of how brains encode and manipulate self- and other- directed meta-cognitive information. (Could Barnden's *ATT-Meta* system be a start?)

Can we get clues from biological evolution?

If a bird is seen to be flying around in an elliptical orbit, it will not be because the bird's motion is caused by elliptical physical motion of something outside the bird, as rotary motion of a leaf in a river whirlpool is caused by the motion of the water. The bird will have information about its environment (e.g. about possible prey, a possible nesting site, or a predator approaching its nestlings) and identified needs (e.g. to get food, to find a good place for a nest, to find nesting material, to distract a predator, etc.) It will also need the ability to increase or decrease speed and change direction. Depending on the circumstances, the bird's motion will use energy (either in its muscles, or in wind or updrafts, or gravity), controlled on the basis of constantly changing information, to produce motion with intended results. There may or may not be additional meta-cognition (self-awareness). Instead of being moved solely by external physical forces, as planets and

[4] Revised edition online at http://www.cs.bham.ac.uk/research/projects/cogaff/crp#chap8

[5] For more detail see http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html

[6] However, it's a fashionable mistake in some circles to assume that mathematical discovery necessarily requires social uses of language, just as its a fashionable mistake to assume physical embodiment plays a role in all mathematical reasoning.

clouds are, the bird has information-processing mechanisms that control its motion. E.g. it can select some information items rather than others then select and execute an action, then switch to a different goal and different action. Evolution changes the amount and variety of information that can be acquired, manipulated, stored and used, and the variety of types of needs and goals that can drive such processes.

Long before humans existed, various mathematical structures and relationships, some but not all numerical, were involved in control processes, including increasing or decreasing turn angles, speed, height, joint angles, forces applied, etc. At some stage humans developed meta-informational (meta-cognitive) abilities to reflect on, reason about, increasingly complex examples of such structures and relationships, including possible future structures; e.g. shelters not yet built, clothing not yet made from an animal skin, a meal whose ingredients are not yet assembled.

As yet unknown evolutionary changes must have supported new proto-mathematical abilities for manipulating and using information about structures, processes, actions, forces, etc. including future possible (intended) cases. A large subset is shared with other intelligent species. The mathematics that we teach and do research on is just a small subset, and almost certainly cannot be understood independently of the less obvious mathematical competences we share with many other species, especially topological and geometrical competences. Different mathematical structures occur in percepts, in intentions, in plans and, later on, in linguistic communications.

Evolutionary pressures for mathematical minds

Increasingly complex forms of life need to use increasingly complex and varied information structures including motives: information states concerned not with what *is* the case but with what *should be* the case, i.e. not just belief-like but also desire-like information contents of increasingly complex kinds. I am not claiming that ALL intelligent behaviour is based on current biological needs, or expected rewards, since some motives are triggered as "internal reflexes" by opportunities without any expected benefit, as can be seen in much playful activity in young children, kittens, apes, and others. What is learnt in such contexts can have consequences that are later useful in ways that the individual could not possibly predict. So although the mechanisms do not involve *expected* rewards, the indirect benefits they previously produced in ancestors may explain the survival of the goal generating mechanisms in their descendants, though not how they formed in the first place (using specially evolved construction kits).

There may be "branch points" during development where different lineages take different branches, under control of genome and environment. But at later stages of development evolution can support greater environmental variation, so that genetically programmed developmental choices may use information previously acquired during development. The fact that common gene-based language potential can support development of thousands of different languages in different contexts illustrates this.

That requires the genome to have a mathematically abstract language specification with very rich generative power, as Chomsky pointed out long ago. I suggest that that is a common feature of biological intelligence, which began with evolution of intelligent control systems in many species that have never been able to use human languages. But they must have rich internal languages for specifying percepts, goals, actions, and environmental structures, including structures that were never encountered by earlier members of the species. A special case is ability to represent entirely new affordances–not unique to humans. .

One of the deep discoveries of evolution was the need for reflexes: actions triggered without the agent having any idea what the benefits are. We need to generalise this to include reflex triggering of new internal motivational states that join other current motivational states, and may or may not lead to action, depending on what else is going on. I call this "Architecture Based Motivation" (ABM) in contrast with "Reward Based Motivation" (RBM) which requires every selected motive to be associated with some measurable expected utility.

ABM seems to be the basis of much exploratory and playful behaviour, including developing linguistic abilities of different sorts, e.g. early babbling and later uses of increasingly complex syntactic forms and growing vocabularies. This may be a source of mathematical development and discovery in young humans (with much individual variation). It also depends on prior, presumably genome-derived, mathematical competences required for exploring novel semantic contents.

As evolution produced increasingly complex organisms, with increasingly complex time-varying needs, and complex articulated bodies capable of rich and varied interactions with the environment, the requirements for mathematical abstraction in information processing increased, including use of geometrical and topological information about spatial structures and both observed and desired changes in spatial relationships, unlike organisms that simply depend on physical influences such as wind or water or the intervention of other organisms to produce the changes they need, e.g. use of other organisms for seed dispersal.

Simple types of information-based control are *online*: information is used as it is acquired and immediately overwritten by new information, e.g. if an animal moves continuously towards a fixed or moving edible target. More sophisticated organisms combine information fragments acquired at different times to produce richer information-structures concerning the environment, e.g. a human (or urban animal) storing and integrating information about the layout of a town and later using the information to work out a route that will reach a new target. This uses *offline* information processing, and *offline* control: actions may be selected long before they are performed, unlike *online* homeostatic control. The richer the environment, the more varied its structures, routes, materials, and other resources, the more powerful the organism's mathematical resources will need be to be able to create and reason about novel possibilities for achieving goals, avoiding dangers, etc. Because of the need to cope with novelty by getting things right first time, empirical learning from repeated trials will be of limited use. This is where mathematical competences are so biologically useful: solutions can be evaluated in advance by reasoning, using structural relationships, instead of having to be evaluated only by repeated testing.

There are differences between a planetary system in which mathematical relationships restrict motions resulting from forces and what goes on in the majority of biological control systems: where, instead of physical processes directly producing or modifying behaviour, there are intervening information processing mechanisms. E.g. sensory systems acquire information and motor control mechanisms use that information in selecting between control alternatives. The control actions may be influenced by information from several sources: e.g. information about an internal need (e.g. for energy-rich food, or for water) can be combined with information about opportunities and obstacles in the environment, or lurking predators. These are unlike processes combining physical forces.

In many cases physical attractive forces increase as distance is diminished, which in the case of physical control leads to increasing acceleration. That could be disastrous for an organism approaching a target: so it is useful be able to detect closeness to the target and use that information to produce deceleration (using stored energy for braking). Where the target is a prey animal that is likely to attempt escape, acceleration right up to contact may be useful, but that requires additional control mechanisms, e.g. producing appropriate motion of claws, or beak or jaws, to capture (or perhaps kill) the prey while avoiding a dangerous impact for the predator.

Even in a very simple single-celled organism, mathematical relationships play a role in control of osmotic pressure, which can be altered by absorption of nutrients or secretion of waste products. One of the important differences between forces and information contents is that forces remain active in the presence of other forces, and their effects combine to produce "resultant" forces, whereas an information item can be temporarily disabled by being ignored, until some urgent task has been completed. So it is essential in organisms to be able to use information to control which other information items have causal powers at which times.

Mathematical competences required for use of such information in selecting and controlling actions are found in many non-human species. These are important aspects of perception and use of what James Gibson called "affordances" in the environment . However Gibson focused on a subset of affordances, mainly those that are relevant to *online* control of actions by the perceiver, whereas humans can perceive and make use of positive and negative affordances for other individuals, and "proto-affordances" – that involve possibilities for change in many aspects of the environment that are not produced by the perceiver and which may be irrelevant to the needs of the perceiver, for example, perceiving that if a certain apple drops off the tree it will not hit the ground because it will land on a rock, whereas if the rock is moved the result will be different. Humans, (and some other organisms?) can also deal with negative affordances that are impossibilities.[7]

Moreover, control relationships can change as an animal grows: genetic mechanisms must somehow enable controlling forces to be varied as sizes, weights, moments of inertia, geometrical relationships and muscular strength change in a growing animal, as D'Arcy Thompson and others have noted.

Besides control based on quantitative relationships, evolution also uses information about *structures* and *structural relationships*, insofar as genetic information plays a role in specifying parts and relationships between parts of developing chemical and physical structures. In humans, another kind of mathematical power is involved in the ability of individuals to develop linguistic competences that make use of complex and varied grammatical structures for information-bearing utterances, and competences that build complex semantic interpretations based on structural relationships (compositional semantics).

Evolution: the blind mathematician

In all these cases the evolutionary and developmental control mechanisms seem to make use of repeated discovery of new structures that can be abstracted from particular instances and later combined with different information in new contexts, while performing complex controlled actions, and while interpreting complex structured perceptual input. Some information about newly discovered abstractions is somehow encoded in genetic mechanisms that allow the information gained to be used in later products of evolution. And in many cases it is crucial that the replication is not a matter of repeated blind copying of the same structure: what is passed on is at a level of abstraction that can be instantiated in different instances, for example (a) when used for continued control of organisms or parts of organisms while growth produces different sizes, weights, size-ratios, moments of inertia, etc. during development and (b) when used in newly emerging species with different details caused by changes in other parts of the genome.

So evolution can be described as "discovering" that new mathematical structures are possible, and that they can be used for new control functions, during reproduction, during development, and during particular actions. Moreover, the evolutionary and developmental histories can be regarded as *proofs* of those mathematical possibilities, even though there is no mathematical mind at work in discovering the theorems or creating the proofs. In that sense biological evolution can be regarded as a "blind theorem prover", rather than Dawkins' "blind watch-maker".

Computers are much faster and more accurate than humans at performing certain kinds of mathematical operations, including numerical and statistical operations, and using arith-

[7] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html

metic, algebra and logic to derive conclusions, solve problems and make plans. But not all mathematical discoveries made by humans are based on arithmetic, algebra and logic. Examples include the ancient geometrical and topological discoveries leading up to Euclid's *Elements*[8] made by ancient mathematicians, e.g. Euclid, Archimedes, Pythagoras, Zeno and others; and also the implicit mathematical discoveries regarding syntactic and semantic structures used in human languages.

Even pre-verbal toddlers, and other animals, such as crows, elephants, weaver-birds and squirrels, seem to have spatial (e.g. topological) reasoning competences unmatched by current automated theorem provers and highly trained robots.[9] However, non-human mathematical reasoners and very young humans lack meta-cognitive abilities to reflect on their mathematical discoveries or to explain and defend them against criticism. That limitation may also have afflicted our adult ancestors who first started to make unreflective and unsystematic use of some of their practical reasoning abilities.

I suspect the variety of evolved mathematical competences is far larger and deeper than anyone has noticed. Researchers are currently struggling to sort them out. E.g. there is a notion of *density* (of grains of salt or sand, of leaves, or flocking birds) and a notion of an area or volume occupied with uniform density, which leads to a notion of amount or numerosity that varies both in proportion with the density and with the area or volume, because *total* amount, or numerosity, as opposed to (cardinal) number increases or decreases as either the density, or the area/volume increases or decreases. Understanding that can lead to inferences about increasing numerosity as density remains constant and area or volume increases, or as density increases while area or volume remains constant. This can support judgements of partial orderings of amount or numerosity. But it does not provide a basis for comparing two regions A and B where area or volume of A is greater than that of B, but density of occupancy of B is greater than that of A. Understanding the tradeoff between change in total space and change in density requires a kind of mathematical sophistication that is a pre-cursor to the understanding of integral calculus. I don't know whether anyone understands the mechanisms used in such cases, nor how they produce new competences during development.

Still more mechanism is required for comparisons of areas, volumes, lengths and amounts of stuff occupying areas or volumes. Those require understanding of new kinds of number that occupy spaces between the natural numbers. Ratios, or fractions may seem at first to suffice, so that we can talk of a jug being half, three quarters, five sixths, full etc., but ancient mathematicians discovered (to their horror) that those ratios do not suffice. In particular something more is needed if the side of a square and its diagonal are to be thought of both having definite lengths, as was understood by the time of Euclid's *Elements*.

[8] http://www.gutenberg.org/ebooks/21076

[9] Examples involving human toddlers can be found here http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html

Limited mathematical abilities of AI systems

Computers are generally thought of as good at doing mathematics. But that is based on a limited view of the scope of mathematics. Computers can perform logical, arithmetical (and therefore statistical) calculations, and operations on text strings, at enormous speeds, because those processes are readily mapped onto operations on bit patterns – especially in combination with random access memory (RAM) operations that allow contents of memory locations to be checked or modified at very high speed (unlike operations on the tape of a Turing machine). Moreover developments in AI, software engineering, theoretical computer science, networking technology, and increasingly sophisticated fabrication processes have expanded the abilities of (networks of) computers so that they increasingly form interfaces to a host of everyday functions, and outperform humans in many activities.

Yet there are many aspects of human (and non-human) intelligence that are not yet modelled on computers, and seem to be particularly hard to model. Many cases go unnoticed by researchers because because they involve not just abilities to act (e.g. catching, throwing, assembling, stacking, etc.) but also abilities to understand possibilities, necessities and impossibilities, which abound in both mathematics and everyday life. These aspects of human and animal intelligence cannot be derived from statistics based learning, nor expressed in probabilistic frameworks, because they are concerned with what is possible, impossible, or necessarily the case, not probabilities. And many are about structures, not measures.

The Turing-inspired Meta-Morphogenesis project includes trying to understand many intermediate forms of information processing between the very simplest organisms and current highly intelligent animals, in the hope that we may stumble across cases that we have never previously thought of that provide new clues regarding mechanisms required and used in brains. The project has already identified a need for evolution to make use of both the fundamental construction kit (FCK) provided by the physical universe and also many derived construction kits (DCKs) produced by biological evolution. Some are concrete construction kits for producing physical and chemical structures and processes. Others are abstract construction kits for producing information structures and information processing mechanisms. It is hoped that eventually we'll understand the sorts of construction kit required to replicate human mathematical intelligence in machines, so that we'll know how to make a baby Kantian robot that can grow up to make discoveries like Euclid.

For more on the meta-morphogenesis project see: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html

FOR MISSING REFERENCES SEE:
http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-iccm17.pdf