What Sort of Information-Processing Machinery Could Ancient Geometers Have Used?

Aaron Sloman*

School of Computer Science,
The University of Birmingham, UK
a.sloman[AT]cs.bham.ac.uk
http://www.cs.bham.ac.uk/~axs

Abstract. There have been many automated geometry theorem provers since Gelernter's 1964 prover discovered a proof of the 'Pons Asinorum' not previously known to the developer. But they all start with logic-based formulations of Euclid's axioms, postulates, etc. (e.g. using David Hilbert's axiomatization, and assuming the Cartesian coordinate representation of geometry). But that is not how the ancient mathematicians started: that approach to mathematics was not developed until centuries later. What sorts of reasoning machinery could the ancient mathematicians, and other intelligent animals be using for spatial reasoning? *"Diagrams in minds"* perhaps? How did natural selection produce such machinery? Which components are shared with other intelligent species? Does the machinery exist at or before birth in humans and if not how and when does it develop? How do brains implement such machines? Could they be implemented as virtual machines on digital computers, and if not what human engineered "Super Turing" mechanisms could replicate what brains do? How are they specified in a genome? Does Turing's work on chemical morphogenesis published shortly before he died indicate that he was interested in this problem? Will the answers to these questions vindicate Kant's claims about the nature of mathematical knowledge, including his claim that mathematical truths are non-empirical, synthetic and necessary? Perhaps it's time for discussions of consciousness to return to the nature of ancient mathematical consciousness, and related aspects of everyday intelligence, usually ignored in current discussions of consciousness.

Key topics
Ancient geometrical reasoning
Non-human mathematical competences
Development of geometrical and topological reasoning
Evolution of geometric/topological reasoning, in humans and other species
Kant on mathematical knowledge
Are relevant brain mechanisms known?
How can brains represent necessity/impossibility?
Possible super-Turing reasoning machinery.

NOTE: In this version all internal references and URLs are clickable, and coloured blue.

1 Introduction

There have been theories of consciousness that make use of mathematics, e.g. mathematical models of patterns of activity in neural nets, but no theory that I have encountered explains how brains enable great mathematical discoveries to be made, e.g. the discoveries in geometry and topology, made many centuries ago, some of which, in Euclid's *Elements*[1], are still in regular use world-wide by scientists, engineers and architects. There have been AI geometry theorem provers since the 1960s, e.g. [2], but they all start with logic-based formulations of Euclid's axioms, whereas for ancient mathematicians the axioms and postulates were not arbitrarily chosen starting points for logical derivations. Rather Euclid's axioms and postulates summarised previously made *discoveries*. They were selected because other geometrical facts could be derived from them, even if they had originally been discovered independently.

* This is a draft paper submitted to a conference – likely to be revised. **Last updated December 28, 2017**.

For example, Euclid showed how the triangle sum theorem (*Internal angles of a planar triangle must always sum to half a rotation (180 °)*) could be derived using the parallel postulate, but Mary Pardoe discovered, while teaching mathematics to school children (around 1970), that the theorem can be proved without reference to parallel lines as shown in Fig. 1[1] What brain mechanisms allow such discoveries to be made, and understood? As far as I know, there is nothing in current neuroscience that explains such capabilities, and no AI system that can make, or even understand, such discoveries.

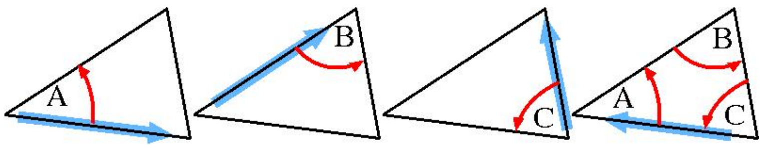

Fig. 1. Mary Pardoe's proof of the triangle sum theorem. Her pupils understood and remembered this more easily than the standard proof, using parallel lines.

There are many geometrical discoveries that can be made without starting from Euclid's axioms. For example, possibilities for creating 3-D structures by repeated folds of a flat sheet of paper (Origami) include cases that cannot be achieved in Euclidean geometry, including trisection of an arbitrary angle.[2]

Moreover, there are various extensions of Euclidean geometry, that are not derivable within Euclidean geometry, including the *Neusis* construction that was known to ancient mathematicians, but not included in Euclid's *Elements*. It involves use of a movable straight edge with two marks, and it also allows arbitrary angles to be trisected easily.[3] The discovery of non-euclidean geometries was another important example, famously used by Einstein in his General Theory of Relativity.

Although Euclidean geometry can be axiomatised using logic and algebra, as David Hilbert showed in 1899 [3], it is clear that the human ability to discover and understand truths of geometry does not depend on use of logical and algebraic reasoning of types that were unknown to ancient mathematicians, and were developed only within the last few centuries.

Topological reasoning ability seems to be even more widespread as discussed in [4]. Young children who have never studied logic or algebra can tell that it is impossible for two linked rings made of solid, impermeable matter to become unlinked without at least one of them changing shape (e.g. ceasing to be a ring). This can be seen in their responses to clever stage magicians who make it *look* as if the impossible has been achieved. A closely related topological problem: if a length of string is passed through a ring, it can be removed from the ring by pulling either end of the string, while the ring is held fixed. Try persuading a child that the string will be removed twice as fast if both ends are held and pulled together. What brain mechanisms enable us to see that such things are impossible? It is not obvious what we would have to add to current AI systems to give them such abilities. Probably a current learning machine would be fooled (at least temporarily) by the evidence that usually two people pulling something manage to get it moving faster than one pulling alone.

I have found some admirers of deep learning mechanisms who believe that given appropriate training such a mechanism could make the same discoveries as ancient mathematicians, and human toddlers can. But those mechanisms are inherently statistics based and can only

[1] For discussion see http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html
[2] https://en.wikipedia.org/wiki/Origami
[3] http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html

discover that certain generalisations have high, or low, probabilities. They cannot discover that something is *necessarily* true or that something is *impossible*: these are totally different from very high and very low probabilities – as Kant understood when he pointed out [5,6] that Hume's classification of types of knowledge was incomplete. Trainable neural nets cannot even *express the idea* of something being impossible, or necessarily the case: without the expressive power, they cannot have the reasoning power.

Kant argued that there are important types of mathematical knowledge that are *non-empirical* and are about *necessary* truths and *impossibilities* (necessary falsehoods), for which statistical evidence can never suffice.[4] Moreover, it is not at all clear what enables humans to understand these concepts: neural nets cannot express necessity or impossibility. That implies that human brains have mechanisms with powers beyond those of artificial neural nets, since humans (and perhaps some other animals) can understand and use those concepts, e.g. in recognizing that performing a certain action is a *guaranteed* way of achieving some goal, or in recognizing that *no action* could possibly achieve the goal.

How can you decide whether there is a spatial configuration in which a planar triangle and a circle have have exactly seven boundary points in common? You can work out which numbers of common points are possible, by doing mental experiments with imagined triangles and circles. As far as I know no current AI system can discover that impossibility without using something like Hilbert's axiomatisation of geometry, using cartesian coordinates, which children don't need, and ancient mathematicians did not know about: Cartesian coordinates were not discovered until the 17th Century. (The discovery was crucial to Newton's mechanics and the invention/discovery of differential and integral calculus, on which a great deal of modern science and engineering depends.)

2 Why Is Non-empirical Knowledge of Non-contingent Truths Important?

The kind of mathematical knowledge under discussion, knowledge of general facts that can be identified without requiring exhaustive testing in large numbers of situations, is not just a philosophical oddity. It is of great practical importance to intelligent agents. For example, knowing that something is impossible makes it unnecessary to waste time to find out whether it can occur. Likewise knowing that having some feature is a necessary consequence of having some other feature allows decisions to be taken and used with confidence that might otherwise have to be tested repeatedly and with caution, wasting time and energy.

For example, given that visual information travels in straight lines except in special situations (e.g. crossing water/air boundaries) and you are looking at a partly obstructed portion of the environment which way should you move to see beyond the obstruction? In many cases the choice is whether to move your head more to the left or to the right, to see round an obstruction with a vertical edge on the left or on the right. I suspect far more animals have that sort of mathematical competence than anyone has ever investigated, even if most don't know they have it or why it works. In those cases, the mathematical solution to the problem may have been directly programmed by evolution.

In the much less common situation of Fig. 2, if you know that the object you are seeking is nearly as tall as the obstructing wall, and it is either infeasible, or unhelpful to move left or right (e.g. the wall is too long), then you can work out that walking *away* from the wall will raise your eye height to well above the height of the wall, whereas moving toward the

[4] However, modal operators, e.g. "necessary", "impossible" should not be analysed using "possible world" semantics. Details are beyond the scope of this paper [7,8].

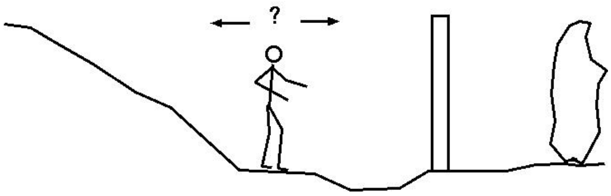

Fig. 2. Should you walk towards or away from the wall to see what is beyond it?

wall will not. If you already know that the ground around you is firm, then you can work out which direction of motion must give you a view over the top of the wall – you must walk uphill to increase your head height. Of course the (mathematical) "must" here is relative to many assumptions about absence of malevolent influences, invisible step ladders etc. There are many practical problems, including engineering problems, where the mathematically best answer is the one to use even if there is no absolute guarantee that the rest of the environment will remain normal.

I suggest that it is not only humans that can benefit from this kind of reasoning. Biological evolution has clearly made and used many mathematical discoveries in selecting both physical or chemical structures and in selecting control mechanisms for those structures. For example, negative feedback control is used in many "homeostatic" control mechanisms from the very simplest organisms to control of blood pressure, temperature, chemical balances and other features of complex organisms. A more complex use of mathematical abstraction is the design of muscular control mechanisms for organisms that vary in size, shape, weights, and moments of inertia of various body parts as they grow. The muscular forces exerted in a child need to be continually replaced by larger forces as that child grows bigger, heavier, and faster-moving. So a parametrised design, with changeable parameters is needed. This is a simple example from a much more complex extended investigation of evolution's discovery and use of construction kits of many kinds, including concrete and abstract construction kits. For more on evolved construction kits and their mathematical properties see[9].[5]

3 Toddler Theorems and Animal Intelligence

In [10] I have collected examples of types of proto-mathematical spatial reasoning and discovery that can be observed in young children without any mathematical training, including pre-verbal children. I have labelled some of the examples "toddler theorems", mentioned below. Piaget recorded many more examples, e.g. in his last two books [11,12] as well as earlier work.

There is also evidence for what could be called "proto-mathematical" reasoning about spatial structures in other intelligent species, e.g. squirrels, the planning and plan execution abilities of Portia spiders [13], and the abilities of nest-building birds [14]. The variety of different spatial configurations that can arise during assembly of a nest from twigs and leaves is so huge that if birds had to learn from experience which configurations are useful which not, and which actions produce the useful configurations in various situations, very few might live long enough learn how to build even one nest reliably. Yet weaver birds manage to produce very complex knotted structures and it is not plausible to suggest that evolution has evolved innate reflex responses for every intermediate situation that can occur during the process of construction of a nest using hundreds or thousands of leaves. Instead, it seems to have provided a kind of implicit mathematical reasoning ability that allows the birds to

[5] Continuing work on evolved construction kits is online here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html

choose between good and bad options in a very wide range of situations. For an online video showing some parts of the construction process see https://www.youtube.com/watch?v=qbWM1QAVGzs.

I am not suggesting that the birds have the kind of mathematical understanding of what they are doing that an engineer does who uses mathematical reasoning to produce novel designs that are provably effective. That requires at least two distinct levels of competence: (a) the ability to reason mathematically in particular cases, and (b) meta- knowledge about how and why that form of mathematical reasoning (e.g. using physical or imagined diagrams) works. In very young humans, and in many intelligent non-human species, I suggest there are evolved spatial-mathematical abilities that allow information about appropriate actions to be derived from a range of intention/situation combinations without any (meta-cognitive) understanding of why those decisions work, or how to extend them to novel situations.

There may also be some (slightly?) more intelligent species that have mathematical (e.g. geometrical, topological) reasoning abilities that allow them to solve novel problems without wasting time on trial and error learning, but without knowing what they are doing or why it works – a description that also fits electronic calculators and many other useful software tools.

Young pre-verbal humans seem to have that kind of unwitting mathematical competence. Several examples of "toddler theorems" are presented in [10], including reasoning about how to avoid jamming fingers when pushing drawers shut. A more complete investigation might reveal several layers of mathematical and meta-mathematical development in young humans, combining genetic factors with information gained from the environment, e.g. the materials in the environment, including types of furniture, types of toy, and types of games played. I think I learnt a great deal from meccano sets. (Compare the processes of "representational redescription" postulated by Karmiloff-Smith [15].)

A particular example of non-spatial mathematical intelligence in young humans is the ability to create subsuming generative grammars after many patterns of verbal communication have been found to work in the environment. This has the great benefit of allowing novel linguistic structures to be created, or to be understood, without first learning how they work from examples.

This stage is usually followed by a further level of competence in adjusting the mechanisms used to cope with exceptions to the rules in the child's linguistic input and output mechanisms.

That is a rather messy kind of mathematical process, and the use of such abilities to derive a new linguistic utterance to communicate a novel thought is not guaranteed to be successful because of its dependence on the competences and vagaries of other humans. (This could be seen as straying beyond the theme of diagrams, though diagrams are often used in the teaching of linguistics to explain which syntactic and semantic structures are in use. I suspect very little is known about the precise forms of representation used in young human brains during generation and comprehension of language, whether spoken, written or signed.)

4 Meta-level competences

A superb, highly original, creative engineer requires at least one additional layer of competence: meta-meta- knowledge about how to search a space of mathematical structures to find a new mathematical technique when faced with a novel problem, in addition to knowing how to test and evaluate particular techniques, and how to deploy the techniques in a variety of situations.

The kinds of mathematical competence required of sophisticated 21st century engineers involve many fields of mathematics that are relatively recent discoveries, including, for example, knowledge of algebra, differential equations, formal grammars, probability theory, the theory of games and decisions, theories of algorithms and data-structures developed in computer science and many more. I am not claiming that biological evolution produced built in knowledge of all those types.

However, in at least some humans, evolution (aided by cultural developments) seems to have produced abilities to absorb hard-won mathematical discoveries that have been made by previous generations, and then use those as a platform on which to build yet more kinds of mathematics, either as a kind of playful activity that is enjoyed for its own sake, or as a goal directed activity seeking a kind of mathematics that allows new solutions to be found for old or recently encountered practical problems (e.g. seeking new mathematics for use in fundamental physics, or linguistics, or AI).

This process depends on a feature of the human genome originally proposed in collaboration with Jackie Chappell, in [16], which I now call "The meta-configured genome".[6] It allows recently developed abstractions from previously evolved competences to be instantiated in novel ways in each generation, possibly building on fairly recent discoveries by previous generations. See Fig 3. This allows greater developmental leaps across generations than could be achieved by use of a fixed learning mechanism provided by the genome. The diagram gives a rough indication of the mechanism in which staggered waves of gene-expression build on and extend previous products of the mechanism, during development of a single individual.

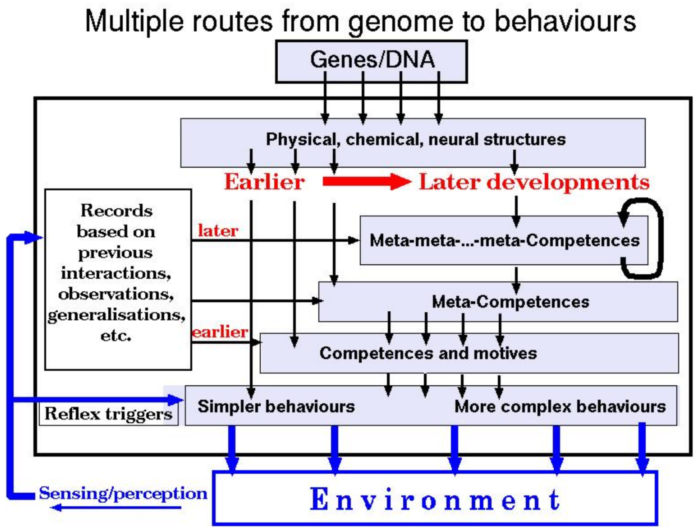

Fig. 3. Staggered "waves of expression" of the Meta-Configured Genome. Layers at the bottom begin development earliest. Processes further to the right occur later, building on records of earlier processes that help to instantiate more recently evolved genetic abstractions, expressed later in development. Many important motives are not reward-based but triggered by powerful internal reflexes produced by a combination of evolution and results of previous development.[17]

The *spatial* reasoning capabilities produced by these mechanisms in humans seem to be very different from the capabilities of mechanisms currently available in AI, including both logic-based reasoning mechanisms (argued by McCarthy and Hayes to be adequate for intelligent systems in 1969 [18]) and the currently more fashionable "brain-inspired"

[6] Partly specified in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html

mechanisms based on statistical learning mechanisms implemented in neural nets surveyed by Schmidhuber in [19].

Neural models of learning, reasoning, planning, and decision making currently used in AI, including robotics, deal with networks of nodes with numerical attributes and linked numerical relationships, whereas the forms of information-processing involved in various kinds of mathematical discovery that make it unnecessary to collect statistical data from samples, are more like spatial manipulation of spatial structures: "diagrams in the mind"[20]. Perception and use of spatial affordances, by humans and other animals acting in natural environments, require abilities to perceive and reason about spatial structures and spatial relationships, including topological relationships such as containment and overlap, and partial orderings (nearer, wider, more curved, etc), rather than precise measures. This idea is developed with illustrative videos and simple examples in [21].

5 Back to ancient mathematical reasoning and discovery

Neither the logic based AI mechanisms recommended in [18]) nor the currently more fashionable allegedly "brain-inspired", but quite un-brainlike neural net mechanisms (which ignore all the chemical complexity of brains), seem to be able to replicate the kinds of mathematical reasoning that led to the deep, ancient discoveries assembled by Euclid.

By examining examples of the spatial (diagrammatic) reasoning involved in ancient mathematical discoveries we may hope to gain some insights into what is missing from current forms of computation. I'll use an example that as far as I know has never been deemed worthy of note by mathematicians, but has a number of interesting features, including very easy comprehension by non-mathematicians.

This is an artificial example, but similar points could be made about various stages of nest construction by birds, though details would be very different depending on the materials and construction processes used. For example, fetching lumps of mud and pressing them onto a surface where the new nest is being built, fetching twigs and weaving them into a stable structure on a tree branch as crows and magpies do, and fetching leaves and weaving them into hanging nests, as weaver birds do[7] all require a collection of abilities to perceive structures, select items to manipulate, moving them to new required locations, and then taking actions to enable the new items to be part of a growing stable structure able to provide support and shelter.

A conjecture: Information processing mechanisms required for practical purposes in structured environments evolved in many species, using geometric and topological reasoning about spatial structures and relationships, without precise metrical information. In humans, those mechanisms were later used in new ways, in conjunction with new meta-cognitive and meta-meta-cognitive mechanisms, that eventually made possible explicit mathematical reasoning, discussion, and teaching, especially reasoning about topological and geometrical aspects of structures and processes in the environment.

It is often assumed that mathematical discovery and reasoning must be concerned with *numerical* values and relationships, but I suggest those came much later and in many cases are not needed because qualitative relationships, such as partial orderings, suffice and are more accessible to biological mechanisms, and adequate for many practical purposes.

In particular, as organisms evolve to cope with more complex structures and processes *in the environment*, they use increasingly complex abilities to create and manipulate new *internal* information structures, representing parts and relationships of external structures and

[7] Illustrated by the BBC here https://www.youtube.com/watch?v=6svAIgEnFvw

processes, and supporting reasoning about consequences of possible actions, as hypothesised by Craik in 1943 [22].

Initially those mechanisms and information structures were used for practical decision making and action control in many species, and in pre-verbal human toddlers, e.g. controlling grasping actions and controlling motion towards desired objects, including avoiding obstacles where necessary.

Later on, newly evolved meta-cognitive mechanisms, for reflecting on and comparing successes and failures of such reasoning processes, allowed new, mathematical, aspects of the structures and relationships to be discovered, thought about, and, in some cultures, communicated and used in explicit teaching and discussion. Much later, via social and cultural processes for which I suspect historical records are not available, the materials came to be organised systematically, recorded in various external "documents", such as Euclid's *Elements* and taught in specialised sub-communities.

6 Towards a super-turing geometric reasoner: deforming triangles

Why talk about deforming triangles? Because I think there are deep, largely unnoticed, aspects of the ways human and non-human animal minds work that are closely connected with the mechanisms underlying important non-numerical mathematical discoveries by ancient mathematicians, i.e. topological and geometrical discoveries. It is not always remembered that for ancient mathematicians the axioms and postulates in Euclidean geometry were not arbitrarily chosen starting formulae from which conclusions could be derived using pure logic: the ancient axioms were all major *discoveries*, using mechanisms still available to us.

Consider mechanisms involved in thinking about what happens to angles of a triangle as it gets stretched by motion of one vertex relative to the other two. I suggest those mechanisms were available to ancient mathematicians, whether they thought of this example or not.

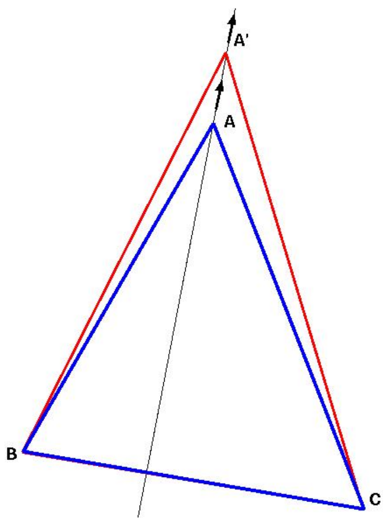

Fig. 4. How does angle A change as it moves further from BC, along a line that passes between B and C, e.g. to A′?

Imagine an arbitrary planar triangle ABC, such as the triangle depicted in Fig.4. What will happen to the angle at A if it continually moves further from the opposite side, BC, along a line that intersects BC and passes through A, as illustrated in Fig. 4. I have informally asked about 40 people this question, all of them academics, from several disciplines, including psychology, philosophy, computing, and AI. All of them seem to find the answer very obvious, with very little thought. Namely as the point A moves further from BC the angle BAC will steadily decrease.

Despite being so obvious to non-mathematicians, this answer has surprising mathematical sophistication. First of all it involves two continua: there is the continuum of locations of the angle A, along the line – or distances of A from the line BC, and the continuum of sizes for the angle A. Second there is a systematic relationship between the two continua: as the distance increases the angle size decreases. This is a qualitative relationship that holds for a wide variety of shapes and sizes of initial triangle, since no units of measurement for the length or the angle are specified, and the initial shape and size are not restricted by the question, especially if posed without a drawn triangle.

It is not obvious *exactly* how the angle size and the length are related, though it is obvious that as one increases the other decreases. This is not true if the line along which A moves does not intersect the line BC between B and C.[8]

There are surprising complexities in the case where the line of motion intersects the line BC outside the triangle, discussed in these two web pages:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html
For now, however, I merely want to raise the question: what sort of reasoning mechanism would enable a future robot to find the answer in the simple case (Fig. 4) as obvious as (adult) humans seem to.

A robot designed by a mathematician might answer the question by using trigonometry to compute a formula linking the angle size to the distance of A from a point on BC, e.g. the midpoint. But the question whose answer all my respondents found obvious was independent of the intersection point, because it was a qualitative answer: as the distance between the point A and the line BC increases, the size of the angle BAC decreases. What sort of spatial reasoning system could we give a future robot that would enable it to find the answer as obvious as humans seem to, and independently of the precise point at which the line of motion of A intersects BC.?

We can ask this question about the required mechanism for all aspects of human spatial reasoning, mathematical and non-mathematical. I think the answers will be different for professional mathematicians who already have various previously known relevant geometric or trigonometric formulae readily at hand. But I am more interested in the non-mathematicians: what kind of mechanism enables them to find the answer above so obvious.

My tentative conjecture is that there is something brains can do that is not explained by current neural net mechanisms nor by current AI models of spatial reasoning using logic or logic plus algebra, trigonometry etc.

We can think of this in terms of how the mechanism might differ from a Turing machine. A Turing machine has a linearly ordered tape, divided into locations each of which can contain exactly one symbol, which could be the ‘empty’ symbol. The symbols allowed have no parts, and therefore have no relationships to one another, such as one being a part of another, or sharing a common part, etc. The Turing machine has a collection of possible *atomic* states (perhaps indicated by numbers) and a machine table that specifies exactly what should happen when the tape head reads one of the permitted symbols in the current location of the tape if it is in the current state.

[8] The case where A moves along a line that intersects BC outside the triangle is discussed in another document, and some of the implications are surprisingly complex, e.g. the intricacies of the problem of finding where the angle size at A is maximal as it moves along a line meeting BC outside the triangle. For details see: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html.

A human thinking about the triangle problem will be doing something different. There will be an imagined state and an imagined change of state. For example the first state could have an angle at A of a certain size and the change of state would be A moving further from BC.

In a Turing machine the states have no internal structure and there is no structural relationship between the symbols that can be on the tape.

In our supposed human-like reasoner there is a change of state, but it is not a discrete change to a uniquely defined next state: rather it is a change specified as a distance increasing (i.e. the point A moves further from BC, but not by any specified amount).

So we need to replace the Turing machine table which has discrete rules, whose conditions and actions are discrete states with something that can reason about a before and after process in which instead of meaningless symbols the machine finds structured changes in a new state: the point A has moved further from BC and as a result the relationship between the two lines meeting at A has changed: the angle is smaller. But that is not an arbitrary rule that has been adopted. Instead the change of size is a necessary consequence of the increasing distance.

How can that necessary consequence be detected by a machine that does not have such a rule explicitly programmed into it? Humans asked about this seem to give different answers. For example one answer is that as A moves further from BC the angles at B and C must "obviously" increase, and therefore the two lines, BA and CA become closer to being parallel and therefore the angle at A must decrease.

A different sort of answer is that as A moves further from BC the two sides AB and AC will have to change direction to point from A towards the old locations for B and C. By comparing the configuration at the original location A and the new location, e.g. A' we can see that the old lines from A must diverge slightly more than the new lines. So the new angle at A' must be smaller.

This suggests a research problem: find a way to specify a type of machine that could replace a Turing machine's tape, tape-head, and symbol table, with something like a membrane on which marks can be made and which can be stretched, rotated, translated, and its new position compared with the old position, to see what has changed.

Is there a minimal set of such transformations from which all the forms of spatial reasoning required for an intelligent animal can be derived?

I suggest that a fruitful Super-Turing machine research project could look at types of transformation that can occur on a retina as a result of various kinds of motion of the perceiver or external objects, and ways of drawing conclusions from such transformations. This will involve replacing the Turing machine tape with some sort of membrane – perhaps more than one, so that "memories" of previous states can be stored and compared with new states allowing conclusions to be drawn about motions of the perceiver and other things in the environment.

This is just the beginning of a still only partially specified research programme, that might give clues as to what sorts of spatial reasoning mechanisms may be implemented in brains and how they can explain the deep ancient discoveries led to Euclidean, then later non-Euclidean geometries being explored. (Fragments of this proposal have been made previously.) Later developments of these ideas will be added here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-geom.html

7 Multi-layered Genome Expression

If epigenetic (gene expression) processes are spread over extended time periods, and later processes are partly influenced by results of interactions between earlier competences and the environment, that allows development to be much influenced by changing combinations of genome expression and cultural expression. Members of such species can then benefit from a deep mixture of products of biologically and culturally evolution, as illustrated by differences in linguistic development in children in different cultures. This can allow powerful and general evolved "construction kits" to be used in novel ways as knowledge and expertise accumulate in members of a species [15,23,16,24].

Very few theories attempting to characterise, explain, or model human consciousness have paid any attention to mathematical consciousness – the kinds of consciousness involved in making mathematical discoveries, such as the discoveries in geometry, topology and arithmetic by ancient mathematicians reported in Euclid's *Elements* [1].

Conjecture: the proto-mathematical mechanisms enabling such mathematical consciousness originally evolved to serve practical requirements of perception, reasoning, planning, and control of actions, in many species, including humans and other animals able to cope with complex spatial control tasks, including avoiding obstacles, climbing rocks and trees, building nests, choosing shelters, obtaining, and eating vegetable matter (e.g. peeling fruit and cracking open nuts), catching and dismembering prey, fighting, caring for helpless young, mating, route-finding, and many more. James Gibson drew attention to some relevant aspects of perception and action [25,26] but his theory of affordances was too narrow and too shallow to accommodate all of these complexities, including the information processing required for coping with geometric and topological complexities – aspects of which had been noticed earlier by Kant [5,6].

Any *complete theory* of consciousness must at least describe, and if possible also explain, using implementable models, aspects of human consciousness involved in those ancient mathematical discoveries, and related aspects of proto-mathematical discovery in very young children and other animals. In that respect all published theories of consciousness, including theories of perception, action, and learning, that I have encountered are mistaken, or at least incomplete.

Neural mechanisms have been proposed to explain mathematical competences: but they omit key features of actual mathematical competences, described above, that neither current AI models, nor current neuroscience theories, seem able to explain. Since my DPhil thesis defending Kant's philosophy of mathematics in 1962, I have been collecting many examples over many years. I began to think about how to use AI to extend and defend Kant's theses after I met Max Clowes around 1969. The problem turned out to be much harder than I realised. Some partial progress reports are presented in online papers, most of them still under development.[9]

The explicit use of these mechanisms for discovery and reasoning about topology and geometry, or the theory of cardinal and ordinal numbers, does not occur in all cultures. So, both aspects of human evolution and also historical contingencies and cultural developments must have influenced the deployment of those mechanisms in discovery, communication, organisation, and formal teaching of ancient mathematics. But I am not claiming that such human activities created the mathematical structures or made the theorems true. In fact, long before

[9] E.g. http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html,
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html,
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html,
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html,
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/cardinal-ordinal-numbers.html

humans made mathematical discoveries, biological evolution made and used mathematical discoveries, for example in design and deployment of homeostatic control mechanisms, using negative feedback to produce or maintain steady states, or to control rates of change.

And even more dramatically, it seems that the evolution of genomes for organisms whose developmental trajectories include changing details of shape, size, forces required, and speeds of actions, required implicit discovery of reusable mathematical abstractions that allowed such genomes to specify designs with fixed structures (e.g. the topology of a vertebrate skeleton, and its musculature) and changing parameters (e.g. the changing sizes, masses, strengths, etc. of parts, during development and the required changes of control parameters for walking, running, jumping, chewing, etc.) Many of these features of evolution had been noticed (e.g. by Thompson [27]) before human engineers had discovered the importance of such "parametric polymorphism" in designs for computer based control systems and other complex multi-functional kinds of virtual machinery, during the last 70 years or so.

Moreover, as Schrödinger pointed out in 1944 [28][10] The molecular structure of polymer-encoded genomes made use of implicit mathematical discoveries concerning use of discrete sequences for encoding information, by biological evolution, including anticipation of some of Shannon's work. This paper is a partial progress report on half a century of Kant-inspired, then AI-inspired research based on many examples of human and non-human spatial competences (examples – not statistical regularities, for good reasons).

Although this project is old, the idea of a Super-Turing Membrane machine is still new and under development. If this paper is accepted for the conference I hope to be able to add more substance. If the associated half day tutorial proposal is accepted, it will go into a lot more detail in many of the problems, in the hope of inspiring some younger researchers to help speed up progress on this project, which aims to address very deep limitations in current AI, neuroscience, and philosophy of consciousness. I suspect Alan Turing may have been working on a problem of this sort when he wrote "The chemical basis of morphogenesis" [29], now his most cited paper. What would he have done if he had not died two years after it was published? I suspect he did not believe the Church-Turing thesis.

Acknowledgments. I am grateful for help and tolerance from many colleagues and students who did not always see the point of my constant harping on these topics, especially as evidence for deep gaps in current AI, psychology, neuroscience, and philosophy of mathematics.

References

1. Euclid, Casey, J.: The First Six Books of the Elements of Euclid. Project Gutenberg, Salt Lake City (Apr 2007) Third Edition, Revised and enlarged. Dublin: Hodges, Figgis, & Co., Grafton-St. London: Longmans, Green, & Co. 1885.
2. Gelernter, H. In Feigenbaum, E.A., Feldman, J., eds.: Computers &Amp; Thought. McGraw-Hill, New York (1964) 134–152 Re-published 1995 (ISBN 0-262-56092-5).
3. Hilbert, D.: The Foundations of Geometry. Project Gutenberg, Salt Lake City (2005) Translated 1902 by E.J. Townsend, from 1899 German edition.
4. Sauvy, J., Sauvy, S.: The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology. Penguin Education, Harmondsworth (1974) Translated from the French by Pam Wells.
5. Kant, I.: Critique of Pure Reason. Macmillan, London (1781) Translated (1929) by Norman Kemp Smith.
6. Kant, I.: Prolegomena to Any Future Metaphysics That Will Be Able to Present Itself as a Science. (1783) Originally Prolegomena zu einer jeden kunstigen Metaphysik, die als Wissenschaft wird auftreten konnen. Several translated versions online.
7. Sloman, A.: Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil Thesis). PhD thesis, Oxford University (1962)

[10] See also http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html

8. Vetter, B.: Recent Work: Modality without Possible Worlds. Analysis **71** (Oct 2011) 742–754
9. Sloman, A.: Construction Kits for Biological Evolution. In Cooper, S.B., Soskova, M.I., eds.: The Incomputable: Journeys Beyond the Turing Barrier. (2017)
10. Sloman, A.: Meta-Morphogenesis and Toddler Theorems: Case Studies (2013) Online discussion note.
11. Piaget, J.: Possibility and Necessity Vol 1. The role of possibility in cognitive development. U. of Minnesota Press, Minneapolis (1981) Tr. by Helga Feider from French in 1987.
12. Piaget, J.: Possibility and Necessity Vol 2. The role of necessity in cognitive development. U. of Minnesota Press, Minneapolis (1983) Tr. by Helga Feider from French in 1987.
13. Tarsitano, M.: Route selection by a jumping spider (Portia labiata) during the locomotory phase of a detour. Animal Behaviour **72, Issue 6** (December 2006) 1437–1442
14. Weir, A.A.S., Chappell, J., Kacelnik, A.: Shaping of hooks in New Caledonian crows. Science **297**(9 August 2002) (2002) 981
15. Karmiloff-Smith, A.: Beyond Modularity: A Developmental Perspective on Cognitive Science. MIT Press, Cambridge, MA (1992)
16. Chappell, J., Sloman, A.: Natural and artificial meta-configured altricial information-processing systems. International Journal of Unconventional Computing **3**(3) (2007) 211–239
17. Sloman, A.: Architecture-Based Motivation vs Reward-Based Motivation. Newsletter on Philosophy and Computers **09**(1) (2009) 10–13
18. McCarthy, J., Hayes, P.: Some philosophical problems from the standpoint of AI. In Meltzer, B., Michie, D., eds.: Machine Intelligence 4. Edinburgh University Press, Edinburgh, Scotland (1969) 463–502 http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html.
19. Schmidhuber, J.: Deep Learning in Neural Networks: An Overview. Technical Report Technical Report IDSIA-03-14 (2014)
20. Sloman, A.: Diagrams in the mind. In Anderson, M., Meyer, B., Olivier, P., eds.: Diagrammatic Representation and Reasoning. Springer-Verlag, Berlin (2002) 7–28 http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#58.
21. Sloman, A.: Predicting Affordance Changes (Alternative ways to deal with uncertainty). Technical Report COSY-DP-0702, School of Computer Science, University of Birmingham, Birmingham, UK (Nov 2007) http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html.
22. Craik, K.: The Nature of Explanation. Cambridge University Press, London, New York (1943)
23. Jablonka, E., Lamb, M.J.: Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life. MIT Press, Cambridge MA (2005)
24. Sloman, A., Chappell, J.: Computational Cognitive Epigenetics (Commentary on [23]). Behavioral and Brain Sciences **30**(4) (2007) 375–6
25. Gibson, J.J.: The Senses Considered as Perceptual Systems. Houghton Mifflin, Boston (1966)
26. Gibson, J.J.: The Ecological Approach to Visual Perception. Houghton Mifflin, Boston, MA (1979)
27. Thompson, D.W.: On Growth and Form. Cambridge University Press, Cambridge (1917) (Revised Edition 1948).
28. Schrödinger, E.: What is life? CUP, Cambridge (1944)
29. Turing, A.M.: The Chemical Basis Of Morphogenesis. Phil. Trans. R. Soc. London B 237 **237** (1952) 37–72