A Super-Turing (Multi) Membrane Machine for Geometers
(Also for toddlers, and other intelligent animals)
PART 2: Towards a specification for mechanisms
(DRAFT: Liable to change)

PART 1: On philosophical background, is available separately at:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-phil.html

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

Any theory of consciousness that does not include and explain ancient forms of mathematical consciousness is seriously deficient.

Parts of a paper on deforming triangles have been moved into this paper.

PHILOSOPHICAL BIOLOGICAL BACKGROUND
In separate document, including:

Different philosophical and scientific goals
Types of (meta-) theory about mathematics
-- Biological justification/explanation
-- Philosophical justification
-- Mathematical arguments supporting claimed discoveries,
-- Philosophical/scientific explanations of how those discoveries are possible
-- Mechanistic explanations of how various processes occur and what they achieve, or fail to achieve.
-- Philosophical/metaphysical ("grounding") explanations
-- Turing's thoughts about intuition vs ingenuity

THIS DOCUMENT
EVOLVED MATHEMATICAL COMPETENCES

- Main focus: understanding and replicating natural information processing
- Varied forms of representation and reasoning in mathematics
- Some logical reasoning is easy for computers: is logic enough?
    Figure Logic: Which inference is valid, and why?
- Discrete and continuous domains
- Non-discrete computation
- Dana Scott on Geometry without points
- Incomplete draft requirements for a Super-Turing Spatial Reasoning (STSR) engine
- Implications for the Super-Turing machine (STM)
- What replaces the Turing Machine table?
- Possible changes to a spatial configuration: The Apollonius problem
    Figure: Points Lines and Circles
- Impossible routes/trajectories
- Figure: Spirals
- Further requirements for intelligent, mathematical perception
- Properties of the Super Turing membrane machine
- Monitoring mechanisms
- Sub-neural chemical computations?
- Is chemistry essential for some animal competences?
- REFERENCES AND LINKS

BACKGROUND

Different philosophical and scientific goals

EVOLVED MATHEMATICAL COMPETENCES

Main focus: understanding and replicating natural information processing

Several goals are presented in the companion paper, which starts "Life is riddled through and through with mathematical structures, mechanisms, competences, and achievements, without which evolution could not have produced the riches it has produced on this planet".

Some of the mathematical structures are instantiated in physical parts of organisms and processes involving them. Others are information structures used in controlling processes of many kinds, including formation of physical structures, uses of body parts, and in the case of humans explicit use of mathematical concepts, and mathematical reasoning. Only humans seem to have the meta-cognitive ability to discover that they can do these things, and to reflect on and discuss the capabilities involved, and teach them to other individuals

The main focus here is on trying to understand the biological information processing mechanisms (the forms of computation, in a generalised sense of "computation") that make it possible for some types of organism (and perhaps future human-made machines) to make such mathematical discoveries and apply them in achieving increasingly complex practical (e.g. engineering and scientific) goals.

But not all the uses of mathematics are conscious or intentional: many are involved in processes of reproduction discussed by Schrödinger(1944). Many aspects of action control, and cognitive development include mathematical processes, including (unwitting) construction and use of grammars. Many biological control functions use negative feedback loops (homeostasis).

Moreover both human language and thought and human visual perception allow structures nested within larger structures in a way that potentially involves indefinite recursion, even though actual cases may have limited depth. In that sense recursion was used long before mathematicians or logicians explicitly recognized its use. (E.g. think of the nursery doggerel: "This is the house that Jack built":
https://en.wikipedia.org/wiki/This_Is_the_House_That_Jack_Built.)

For now I am not concerned with "rational reconstruction" processes, e.g. attempts to specify requirements and techniques for improving mathematical rigour, or attempts to find some minimal subset of mathematics from which the rest can be derived mathematically, or attempts to draw boundaries between mathematical and non-mathematical concepts, knowledge, and forms of reasoning. Instead this project attempts to understand what actually happened, at a high level of abstraction -- i.e. what evolutionary steps, physical environments, and cultural processes, made it possible for evolution eventually to produce minds capable of the great ancient mathematical discoveries.

There is a primitive, and common, form of mathematical discovery (not the only form), that involves noticing a regularity without understanding why that regularity exists and cannot be violated. Pat Hayes once told me he had encountered a conference receptionist who liked to keep all the unclaimed name cards in a rectangular array. However she had discovered that sometimes she

could not do it. She found that frustrating and blamed it on her own lack of intelligence. She had unwittingly discovered empirically that some numbers are prime, but without understanding why some are prime and some not, or what the mathematical implications are.

A child with unrecognized mathematical talents may discover that some collections of cards of the same size and shape can be rearranged on a flat surface to form a rectangular array of rows and columns, with 2 or more rows and 2 or more columns, and, some time later, discover that some collections of cards cannot be reorganised in that way, although every collection can be arranged in a single row or column.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#primes

However, there is a difference between merely noticing a regularity or lack of regularity (like the receptionist) and understanding why the regularity does or does not exist, in this case understanding why not all numbers are prime. This follows obviously from the fact that for any two numbers N1 and N2 (possibly the same number) a regular array of objects can have N1 columns and N2 rows. Equivalently, it is possible to form a pile of N1 objects and then make copies of that pile until there are N2 piles.

A mathematical mind can grasp not merely those facts but also that the sizes of the piles or arrays that are possible are not constrained by what size pile of cards he or she or anyone else can handle, or whether they could fit on this planet, or within this solar system.

We can distinguish (at least) two ways of trying to understand how such a mind is possible:

1. Understanding what sorts of physical mechanisms could be used in the construction of a working mind of that sort, and how they would need to be used.
2. Understanding how biological evolution (and its products, e.g. social systems) can bring it about that some individuals are able to make mathematical discoveries, understand them, use them, teach them, etc.

But understanding how such mathematical minds are possible will involve understanding how various mathematical exploration and discovery processes in evolution are possible.

Whether such minds can exist in this universe will depend on mathematical features of the physical universe that make the required evolutionary processes and the required physical and chemical facts that make the required brains possible.

I suspect that because most physicists do not attend to such questions about how physics makes biology possible (although Schroedinger did in 1944) most work on fundamental physics ignores relevant constraints on required explanatory powers of the physical universe.

For example, it could turn out that the vast networks of numerical relationships between numerical values that characterise modern physics (e.g. as presented by Tegmark, 2014, among others) include deep structural gaps that physicists will not notice until they try to explain more of the fine details of life, including development, reproduction, evolution, and ever increasing sophistication of information processing of organisms, especially mathematically minded organisms.

Varied forms of representation and reasoning in mathematics

Over many centuries (or perhaps many millennia) humans have found means by which they can derive new mathematical (e.g. topological, geometrical and later numerical) truths that they can discuss explicitly and teach their children to use, some of which they can use (centuries later) to design machines that reach some of the same results more quickly and more reliably than humans do. Diagrams, drawn or imagined, played important roles in many of the discoveries.

There are (at least) two different views of ancient mathematical reasoning using diagrams and words. One view regards the diagrams as mere useful aids (dispensable props?) supporting a kind of diagram-free thinking using logical, arithmetical and algebraic forms of representation identified and studied in great depth centuries after the original geometrical discoveries were made.

Most researchers on foundations of mathematics now seem to focus on investigating such formal modes of reasoning, using discrete symbols, based on, or using, forms of reasoning developed by mathematicians in the last few hundred years. That work typically ignores the very different modes of reasoning that must have been used by our ancestors when they first made the discoveries leading to number theory, topology and geometry as we now know them.

The other, older, view (found in 1781 for example) is that diagrammatic forms of reasoning play a crucial role in some ancient and modern forms of reasoning and they are as much parts of mathematical reasoning as numerical, logical and algebraic reasoning developed in the last two centuries.

In 1938, Alan Turing, in his PhD thesis,

Systems of Logic Based on Ordinals in
Proc. London Mathematical Society, pp. 161-228, 1938
https://doi.org/10.1112/plms/s2-45.1.161

made a distinction between mathematical *intuition* and mathematical *ingenuity*, claiming that the latter (ingenuity) but not the former (intuition) could be implemented, or accurately modelled, in computers.

I conjecture that what he wrote indicated that he had rediscovered a variant of Immanuel Kant's philosophy of mathematics 1781, as explained in Turing(1938).
But as far as I know, Turing never presented a detailed specification for a kind of information processing machine that would have (human-like) mathematical intuition. Perhaps he suspected that some of the (sub-neural) chemical processes in brains were required, which might explain why he wrote Turing (1952), which is very different from all his previous work. But the evidence for that interpretation is very thin.

Whatever Turing actually thought, or wrote, I suggest that trying to design a kind of information processing machine (computer) that is at least partly non-digital, i.e. using combinations of discrete and continuous processes interacting with one another, under the gaze of a suitable "interaction recognizer" might lead us into a deeper understanding of how mathematical brains work, including how they acquire and use novel insights.

This is vaguely hinted at by what Turing wrote about mathematical intuition in 1938 discussed in [Turing(1938)](#).
(Some of this paper needs to be revised in the light of that one.)

In part, the adequacy of proposed mechanisms for mathematical purposes will depend on whether such forms of reasoning can meet the requirement for mathematical discovery to produce modal information about what is possible, impossible, and necessarily the case, as opposed to what we expect to find, or find only some of the time.

Merely collecting statistical evidence and deriving probabilities could not explain the ancient mathematical discoveries in geometry and arithmetic, concerning what is *impossible*, or *necessarily* the case.

In other words, mathematical discoveries go beyond mere information about what actually has, or has not, occurred, or the relative frequencies of various occurrences. I.e. they go beyond observed statistical regularities or derived probabilities, on which much (but by no means all) work in AI now focuses.

> **NOTE:**
> Most philosophers of mind and theorists who write about consciousness seem to ignore mathematical consciousness. This is probably because they don't realise (as Kant did) that mathematical discoveries since ancient times are deeply connected with perception of, including consciousness of, spatial structures and processes and abilities to reason about what to do and what can and cannot be done in a physical environment. So mathematical competences are part of the intelligence of pre-verbal humans and many other species including corvids, squirrels, elephants, octopuses, orangutans, ... and many more. So any philosophical, psychological, neural or biological theory of consciousness that ignores mathematical consciousness is seriously deficient. However, the connections are sometimes implicit rather than explicit. The work of [McClelland, 2017](#) on mental affordances is relevant to our goals.
> For an incomplete discussion of "conscious" as a polymorphous concept see the section on consciousness here:
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html

Some logical reasoning is easy for computers: is logic enough?

A lot of work has been done on automated theorem proving using logic, arithmetic and algebra. This reasoning uses manipulation of notations allowing only discrete transitions, e.g. removing a symbol, adding a symbol, adding or removing a premise, substituting a symbol (simple or complex) for a variable, or checking that a finite, discrete sequence of formulae instantiates members of a finite list of inference rules.

The treatment of elementary propositional logic using truth-tables is similarly discrete. Elementary examples are concerned only with sets of possibilities in a finite space of possibilities. For example, in propositional logic, students (at least those taught by me) learn how operators related to the words, "not", "and", "or", "if ... then..." can be defined using truth tables. (I think starting such teaching with axiomatic systems, or "natural deduction" systems instead is educationally seriously misguided.)

In this framework the validity of inferences can be checked using truth tables, though the sizes of the truth tables grow exponentially with the number of propositions (or propositional variables) involved.

A toy example is reasoning in propositional calculus, illustrated here:

Figure Logic: Which inference is valid, and why?

Premises:
  P or Q
  not Q
Conclusion:
  :- P

Premises:
  P or Q or R
  not Q
Conclusion:
  :- P

By starting from the truth tables for "or" and "not" you can reason about which possible combinations of truth values for P and Q (and R in the second example) will make the premises true and ask whether the conclusion can be false in any of those cases.
If it is impossible for all the premises to be true while the conclusion is false, then the inference is valid. One of these is not valid.

In the first (upper) example, where only two propositions P and Q are involved there are only four possible combinations of truth values to be considered. And that makes it easy to discover that no combination makes both premises true while the conclusion is false.

In the second case, for each of those four combinations R may be true or false, so the total number of possibilities is doubled. But it is still a finite discrete set and can be examined exhaustively to see whether it is possible for both premises to be true and the conclusion false. I assume the answer is obvious for anyone looking at this who understands "or" and "not". Checking for validity in propositional calculus involves exhaustive search through finite sets of discrete possibilities, so it is easy to program computers to do this. More generally, the discreteness of the structures and transitions and the fact that there are only finitely many elementary components with finitely many combinations of truth-values, that allows such reasoning to be modelled on Turing machines, and their descendents -- digital computers.

Things get more complex if the propositional variables, e.g. P, Q, etc. instead of being restricted to two possible truth values (T and F) can take other intermediate values, or even vary continuously. The simple discrete reasoning, based on truth-tables, using or and not etc., will have to be replaced by something mathematically much more complex.
  (a topic investigated by Tarski, Zadeh and others -- e.g. in "Fuzzy Logic"
  https://en.wikipedia.org/wiki/Fuzzy_logic)

That option will not be discussed further in this document, as it is not relevant to my main concern -- to understand very ancient geometrical and topological mathematical reasoning, concerning spaces of possibilities that are not composed of discrete combinations of discrete units, but continuously variable angles, lengths areas, curvatures, volumes, etc. Moreover, in a Turing machine or digital computer, elementary components are cleanly separated and it is not possible for one component to gradually occupy more of the space occupied by another, whereas in geometrical reasoning, illustrated below, entities can be superimposed moved continuously, and their shapes, parts and relationships can change non-discretely, leading to a requirement for a form

of mathematical intelligence that can grasp commonalities in infinitely varying structures.

> **NOTE:** This point is also explained in the video presentation for an invited talk at IJCAI 2017, available here:
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html
> The video starts with examples of spatial and topological perception and reasoning by birds and a pre-verbal toddler. The toddler with pencil example is available separately here
> http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/small-pencil-vid.webm

Many more examples of spatial intelligence, involving squirrels, elephants, nest-building birds, and other animals can be found by searching online videos. They are all relevant because the ancient mathematical abilities that are the core topic of this document build on, and are tightly connected with, a wide range of evolved perceptual, control, and decision-making capabilities in pre-verbal humans and other animals, about which I suspect very little is currently understood.

When those abilities to perceive, reason, and plan are properly described it turns out that they are all missing from current AI programs and robots, including many that display impressive actions after much training or very careful programming (e.g. Boston Dynamics robots). For more on the discrepancies, see Sloman (2007-2014). I am not claiming that it is impossible to give future robots the required reasoning abilities, merely that there are important features of those abilities, closely connected with Kant's philosophy of mathematics, that are rarely noticed by robot designers, psychologists or neuroscientists.

The gap to be bridged is between acquiring statistical generalisations from samples of a space of possibilities, and grasping *necessary* connections and *impossibilities* in the space. Necessity and impossibility are not points on a scale of probabilities.

Discrete and continuous domains

For discrete domains such as propositional calculus and the domain of *finite proofs derivable from some finite set of axioms using a finite set of derivation rules*, there are mechanisms that can determine necessity and impossibility by exhaustive analysis of discrete sets of cases, sometimes supplemented by inductive reasoning to deal with unbounded collections of cases. However, if there is no upper bound to the length of proof that may be required for a particular formula, then exhaustive analysis may be impossible, e.g. if the shortest proof that answers a question has more steps than the number of electrons in the universe.

But for domains like Euclidean geometry and its topological basis, variations in size, shape, and relationships are continuous, not discrete. For example, the space of possible locations of the pencil and sheet of paper, and the space of possible trajectories for the pencil is continuous. So also is the space of possible spiral shapes or possible shapes of a polygon as one vertex moves while the rest are fixed, or variations in relationships as one geometric structure rotates relative to another, or as distances between parts of structures vary.

The sets of possibilities generated by those continuous variations are inherently unbounded and therefore cannot be exhaustively examined in order to determine that within some variety of cases a particular relationships will always hold (necessity) or that it cannot hold (impossibility).

That means that reasoning machinery in such domains needs to be able to find discrete subdivisions between subsets of continuously varying classes of cases, in addition to finding impossibilities and necessary connections between features of geometrical structures or processes.

Such abilities were used repeatedly in the ancient diagrammatic modes of reasoning discovered or recorded by Archimedes, Euclid, Zeno, Pythagoras and others. They were also part of my own experience learning (and enjoying) Euclidean geometry at school about 60 years ago.

My 1962 DPhil thesis) was an attempt to expound and defend the ideas about such ancient mathematical modes of reasoning that I encountered in Kant's *Critique of Pure Reason* (1781) as a graduate student, which I hoped could be given a deeper, more precise, justification using AI techniques after Max Clowes introduced me to AI in 1969.
(Compare Sloman 1971, and 1978 Chapter 7.)

For many years I suspected that the required forms of reasoning, and the forms of spatial perception on which they are based, could be implemented in digital computers using suitably designed virtual machinery.

Key features of *virtual machine functionalism* (not understood by most philosophers who discuss computational models of mind) are summarised here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html

In Sloman 1978 I argued that digital computers must suffice for implementing human-like minds, since any continuous process can be simulated as accurately as required by a discrete process. At that stage I think I was unaware of chaotic systems in which arbitrarily small differences can have arbitrarily large consequences very rapidly.

Another important fact is that there are ways of thinking about continuous structures and processes that yield deep insights. For example, one of the assumptions made implicitly by Euclid, but not formulated as an axiom was an assumption of "completeness". For example, if **C** is a closed continuous curve in a plane and **L** a straight line segment in the same plane, and part of L is in the interior of **L** and part not, then there must be a point **P** that is common to **C** and **L**. This might not be true in all possible spaces, (for example, if straight lines are approximated in a rectangular grid then there could be a pair of lines that cross over each other without sharing a point of intersection, e.g. if one line occupies a diagonal collection of points, the other line could cross it without the two lines sharing a common point. But in Euclidean geometry that's not possible: if two lines in the same plane cross over each other then there must be a point common to both lines, in that plane.

Non-discrete computation

Since working on the Turing-inspired Meta-Morphogenesis project (begun around the year of Turing's centenary, 2012) I have started taking seriously the suggestion that we need to explore alternatives to digitally implemented forms of computation, especially as chemistry-based mechanisms are so important for all forms of life and especially in brains. Chemical mechanisms provide a deep blend of discreteness and continuity, demonstrated for example in Turing (1952).

All of this lends support to the conjecture that there are forms of information processing, and especially mathematical discovery, that have not yet been fully understood and may not be easy (and perhaps not even possible) to implement on digital computers.

This paper is an early attempt to specify a (speculative, and still incompletely defined) alternative to digital computers (e.g. Turing machines), that could provide a basis for many hard to explain aspects of animal perception, reasoning, and problem solving, and which could also explain how the deep discoveries made by ancient mathematicians were possible.

This seems to be closely related to the distinction made in 1938 between two types of process in mathematical discovery, those involving *insight* and those involving *ingenuity*. Turing suggested that computers (e.g. Turing machines) can achieve only the latter. That raises the question: what mechanisms allow brains to achieve the former: *insight*. That topic is discussed further in Turing(1938).

> **NB**
> A possible outcome of this investigation could be discovery of discrete mechanisms that **are** able to support virtual machines of the kind discussed here. In that case we would not need to implement the geometric reasoning mechanisms using non-discrete technology (e.g. something similar to sub-neural chemical mechanisms), although that may remain the only way to achieve very fast, physically compact, low energy, implementations.

I have been calling the new type of machine (provisionally) the Super-Turing Membrane machine, or Super-Turing Geometry machine. There have been previous proposals for Super-Turing computing machines, but not, as far as I know, in the context of producing a robot mathematician able to make discoveries in Euclidean geometry.

It is sometimes forgotten that the axioms of Euclidean geometry were not arbitrary assumptions assembled to specify a form system that can be explored using logical inference methods.

Those axioms were all important discoveries, which seem to require mechanisms of reasoning that we don't yet understand. I don't think current neuroscience can explain them and they are not yet included in AI reasoning systems. So a suggested role for the new type of machine is as part of an explanation of the early forms of reasoning and discovery that led to Euclidean geometry, long before there was a logic-based specification using cartesian coordinate representations of geometrical structures and processes.

The examples below are merely illustrative of the possible roles for the previously unnoticed type of machine, for which I cannot yet give a precise and detailed specification. These are still early, tentative, explorations.

Dana Scott on Geometry without points

Deep and challenging examples of the kind of reasoning I am trying to explain are in this superb (but not always easy to follow) lecture by Dana Scott at Edinburgh University in 2014

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

He assembles and develops some old (pre-20th century) ideas (e.g. proposed by Whitehead and others) concerning the possibility of basing Euclidean geometry on a form of topology that does not assume the existence of points, lines, and surfaces, but constructs them from more basic notions: regions in a point-free topology.

Although it may be possible to produce a formal presentation of the ideas using standard logical inferences from axioms, his presentation clearly depends on his ability, and the ability of his audience, to take in non-logical, spatial forms of reasoning, supported diagrams, hand motions, and verbal descriptions of spatial transformations.

Perhaps the ancient geometers could have discovered that mode of presenting geometry, but they did not. It seems to be a fundamental feature of geometry (the study of spatial structures and processes) that there is no correct basis for it: once the domain is understood we can find different "starting subsets" from which everything else is derived.

And sometimes surprises turn up, like the discovery of the Neusis construction that extends Euclidean geometry in such a way that trisection of an arbitrary angle becomes easy, whereas in pure Euclidean geometry it is provably impossible. For more on that see
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html

One way to seek requirements for the missing mechanisms supporting the ancient geometrical insights is to look carefully at examples of uses of spatial perception, and investigate the kinds of reasoning mechanisms that are required to explain them. Some examples are the questions that can be asked about the processes we perceive, and can think about, in video recordings of changing geometric configurations or viewpoints, as illustrated in these videos:
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/chairs

The non-logical (but not illogical!) forms of representation, that seem to be involved in the original geometric and topological discovery processes and much current human reasoning about spatial structures and processes, seem to have used what I called "analogical" rather than "Fregean" forms of representation in Sloman 1971 and in chapter 7 of Sloman 1978 (not to be confused with representations that make use only of isomorphism).

Further examples of this essentially spatial rather than logical kind of reasoning and discovery are presented in a collection of web pages on this site, including, most recently (Nov 2017, onwards):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html
and these older explorations (some of which are available in both html and pdf formats):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/torus.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/p-geometry.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/rubber-bands.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/shirt.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/chairs
    Studying the spaces and views surrounding parts of visible objects.
Sloman (2007-14) (explaining why AI vision and action control mechanisms could benefit from such mechanisms),

and this video presentation (for a workshop at IJCAI 2017, August 2017):
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/ijcai-17/ai-cogsci-bio-sloman.webm

Incomplete draft requirements for a Super-Turing Spatial Reasoning (STSR) engine

A Turing machine uses a linear tape of discrete locations, each of which can contain at most one symbol from a fixed set of possible symbols. The STSR machine does not have a collection of distinct spaces, each occupied by one of a set of discrete "atomic" symbols, ad in a Turing machine.

Instead it has a collection of more general spaces, each which is capable of being occupied by continuously changeable more or less complex structures.

The simplest Super-Turing machine would have only one such space, which in a very primitive organism might be a sort of map of current sensory stimulation and perhaps also located output signals. The structure of the map can change over time, with components moving around, or changing size, orientation, or shape, either under external or internal influences or both.

A more complex ST machine would have multiple such spaces, with different partly related contents. For example there might be spaces recording patterns of stimulation from different sensor systems, or from the same type of sensor located at different parts of the body, e.g. sensing contact, pressure, motion, temperature, etc.

For some of the sensor spaces the machine can create and manipulate copies, in which contents change in various ways, including changing continuously, by sliding, stretching, rotating, etc. -- unlike a Turing Machine tape whose cell-contents can only switch from one object to another, where the internal structure of the object plays no role, only its identity in a collection of available objects.

In the ST machine two or more shapes in the same space can move independently including moving over or through each other. The resulting interactions of their structures can be inspected and recorded for future reference, or transmitted to another part of the brain, or used immediately to trigger some new process (e.g. a blinking reflex caused by detecting sudden rapid approaching motion).

Shapes in this space can change their location, their size, their orientation, and relative speeds of motion. Moreover, two non-overlapping shapes can move so that they overlap, for example causing new shapes to be formed through intersections of structures.

For example, two initially separate lines can move in such a way as to form a structure with a point of intersection and three or four branches from that point. If the point of intersection moves, then relative lengths of parts of the structure will change: e.g. one getting smaller and the other larger.

A single moving line may change so that it crosses itself with different parts of the line being identifiable as occupying space to one side or the other of the crossing location on the line.

The structures and processes are not restricted to points, lines, circles and polygons: arbitrary blobs, including blobs that move and change their shape while they move can occur in the space. If two of them move they can pass through each other, producing continuously changing boundary points and regions of overlap and non-overlap.

Groups of items, e.g. regular or irregular arrays of blobs and other shapes can exist and move producing new shapes and processes when they interact. How they move should correspond to various physical situations: e.g. the visible silhouette of a complex 3-D object may go through complex changes as the orientation of the object to the line of sight changes. Compare hand shadow art, in which interpretations of shadows of hands vary enormously: https://www.youtube.com/watch?v=4drz7pTt0gw.

The visual space in which percepts move is NOT required to have a metric -- partial orderings suffice for many biological functions, and many other purposes, as illustrated in Sloman (2007-14), although very precise metrics are required for some activities, e.g. playing darts, trapeze displays, and death-defying leaps between tree branches used by spider monkeys, squirrels and others. That paper shows why the explanatory relevance of the ideas presented here extends far beyond mathematical competences.

A full specification of the types of shape, the types of change (passive, externally caused change or active internally controlled change), the types of relationship that can be detected, is a long term research project.

NOTE: At a later date it will be necessary to generalise the contents of these internal information stores to include non-spatial items such as smells, tastes, colours, tones, harmonies, etc.

The sorts of mechanism described here seem to be required for many aspects of human and non-human intelligence, including, for example, the ability of a crawling baby to work out how to move to shut a door with his legs after passing through the doorway, as depicted in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#door-closing and also in much ancient mathematical reasoning, along with many applications in painting, architecture, engineering design, etc.

The (still unknown) conjectured implementation machine is used both for many practical activities in which physical actions are considered, selected, intended, performed and controlled, and also used (at a much later stage in biological evolution, followed by cultural evolution) in ancient mathematical discoveries.

At an early stage there may be only sensory-motor structures that relate to current static or changing states. Those are all online information structures. At a later stage in evolution, and in individual development, there will have to be structures that have different offline functions, including records of past states and processes, representations of predicted states and processes, and representations of desired states and processes, that drive processes to reduce discrepancies between what's current and what's desired. (As in the ancient General Problem Solver (GPS) of Simon and Newell, but with much more richly varied information structures.)

The underlying machine seems to have some features partly analogous to data-structures used in some computer-based sketch-pads, either using materials supporting arbitrary resolution (most unlikely!), or else uses to which resolution of images is irrelevant, with additional meta-cognitive

apparatus as suggested below.

For example if you look at a scene, e.g. a long straight road, and imagine how its appearance will change as your viewpoint changes, e.g. if you move vertically upwards while still looking along the road, the bounding edges of the appearance of the road may increase in length and the perceived angle at which they meet in the distance will vary in ways that can be discovered by experimenting with examples. However, there is no need to experiment with actual examples, since this is a type of correspondence between two spatial changes that can be discovered merely by thinking about the process. E.g. suppose you are standing on a straight road with parallel edges and looking along the road at some fixed location on the road. If the road lies in a planar surface and your viewpoint starts moving upward perpendicular to the surface, while you continue to focus on the same distant part of the road, you may, with a little practice be able to work out how the appearance of the left and right edges of the road will change as you move up. This imaginative ability is probably not innate, but neither is it based entirely on experience of moving upward from a position on a road. I suggest that most people can use their understanding of space to *work out* how the appearance will change, just as you can use your understanding of arithmetic to work out the product of two numbers that you have never previously multiplied.

A some stage this project should include a survey of a wide range of special cases of such relationships between view location, view direction, and orientations of surfaces that we all take for granted in everyday life, and which current robots might have to learn through collecting observations as they move around in space, whereas humans can discover such correlations merely by thinking about them (although abilities to do this may vary across individuals, and with age and development within individuals).

That ability to make such discoveries merely by doing spatial reasoning (possibly with your eyes shut) requires use of sophisticated mathematical mechanisms in brains whose operation could be mimicked on computers by using cartesian coordinate representations of space and performing algebraic and trigonometric transformations of spatial information, although it is very unlikely that that is how brains derive the information about effects of change. On the other hand, as far as I know such abilities cannot yet be explained by known physical features of brains.

In particular, these competences use neither statistical correlations between discrete categories (which would not provide the required kind of mathematical necessity), nor exact functional relationships between metric spaces which neural mechanisms don't seem capable of computing.

A machine used to implement these capabilities will need fairly rich non-metrical topological structures and reasoning powers. E.g. thinking about a plane containing a straight (i.e. symmetrical) line passing through the interior of a convex closed curve in the plane, with no straight portions, e.g. a circle or ellipse, we can see that there must be exactly two points where the line crosses the curve, i.e. connects the interior and the exterior of the curve. If there are more than two the curve cannot be convex. (Why?) If the curve is fixed, the line can be continuously moved "sideways" in the shared plane, relative to the curve until it no longer passes through the interior of the curve, though it will remain co-planar with the curve. As the line moves the points of intersection with the curve will move, eventually merge and then no longer exist.

What can you infer about how the distance between the points of intersection changes while the line moves, given that the curve is closed and convex? How do you do it without knowing the exact shape of the curve?

That's a very special case, but I think ordinary visual (and tactile?) perception involves a vast collection of cases with all sorts of partial orderings that can be used for intelligent control of behaviour using abilities to reason about how as one thing increases another must also increase, or must decrease, or may alternate between increasing and decreasing, etc. (E.g. think of a planar trajectory curving towards obstacle A and away from obstacle B then changing curvature to avoid a collision with A). (See also the "changing affordances" document, mentioned above.)

These mechanisms will be connected to spatial perception mechanisms, e.g. visual, tactile, haptic, auditory, and vestibular mechanisms, but the connections will be complex and indirect, often directly linked to uses of perception in controlling action, as opposed to merely contemplating spatial structures and processes.

Some of my ideas about this are based on ideas in Arnold Trehub's book Trehub(1991), which postulates a central, structured, multi-scale, multi-level dynamically changing store of different aspects of spatial structures and processes in the environment. Visual perception will be one of the sources. (As far as I know he did not discuss examples like mine, nor attempt to explain mathematical cognition.) On this view the primary visual cortex (V1) is best viewed as part of a rapid device for sampling information from what Gibson (1979) called "The optic array". Among other things Trehub's mechanism can explain why the blind spot is normally invisible, even though that was not Trehub's intention!

The tendency in Robotics, and AI generally, to use metrical spatial information, rather than multiple, changing, partial ordering relationships with inexact spatial measures, leads to excessive reliance on complex, but unnecessary, probabilistic reasoning, where online control using partial orderings could suffice, as suggested in Sloman(2007-14). But there are many details still to be filled in.

Implications for the Super-Turing machine (STM)

NOTE: *The label "STM" is often used to refer to a "Short term memory" mechanism (or, more appropriately, a variety of types of short term memory mechanism, since there are clearly several kinds). I use it here for "Super Turing Machine" with only minor qualms, because if the conjectures presented here are substantiated that will transform our ideas about functions and mechanisms of short term memory in humans and other intelligent animals.*

Further progress on requirements will require consideration of the following points:
Instead of the TM's discrete, linear tape, the STM has some still unknown number of overlapping stretchable movable transparent membranes on which 2-D structures can (somehow?) be projected and then slid around, stretched, and rotated and compared.
E.g. this could combine visually perceived surface structures
and the same structures perceived using haptic and tactile sensing.

The TM tape reader that can only move one step left or right and recognize one of a fixed set of discrete symbols would have to be replaced in the STM by something much more sophisticated that can discover structures and processes that result from those superimposed translation and deformation operations, and relative motions between contents of different membrane layers. (E.g. 'watching' a triangle move across a circle, with and without stretching and rotation.)

In ways that are still unknown, the machine needs to be able to detect that whereas some consequences of those transformations are contingently related to previous states, in other cases the structural relationships make the consequences inevitable: especially topological consequences such as possibility or impossibility of routes between two locations that don't cross a particular curve in the same surface.

Such a machine should also be able to use still unknown kinds of exhaustive analysis to reach conclusions about the *impossibility* of some process producing a certain kind of result.

In some organisms there would be Meta-Cognitive Layers inspecting those processes and, among other things, noting the differences between possible, impossible, contingent, and inevitable consequences of structural relationships (the "alethic" modalities that are central to mathematical discovery).

In TMs and digital computers, that sort of detection can be done by exhaustive analysis of possible modifications of certain symbols, e.g. truth tables, or chains of logical formulae matched against deduction rules; but it's very hard to see exactly how to generalise such meta-cognitive abilities to detect impossibility or necessity in the STM, where things can vary continuously in two dimensions.

Finally, the machine-table of a TM would have to be replaced by something much more complex that reacts to combinations of detected structures and processes in the STM by selecting what to do next. In addition, the biological STM would have to be linked to sensors and effectors, where some of the sensors can project new structures onto the membranes.

The general design would have to accommodate different architectures. For example, I am pretty certain that vertebrate vision would have started without stereo overlap, as in many birds, and many non-carnivorous mammals.

There would need to be something like left and right STMs and mechanisms for transferring information between them as an organism changes direction and what's visible only in one eye becomes visible in the other eye. The same mechanism could be used with greater precision as evolution pushed eyes in some organisms towards the front, producing partly overlapping projections of scenes, making new kinds of stereo vision possible.

(Unfortunately, Julesz' random dot stereograms have fooled some people into thinking that *only* that low-level "pixel based" mechanism is used for biological stereo vision, whereas humans (and I suspect many other animals) can obviously make good use of larger image structures, e.g. perceived vertical edges of large objects, in achieving stereo vision. https://en.wikipedia.org/wiki/Random_dot_stereogram)

I suspect evolution also eventually discovered the benefits of meta-cognitive layers, without which certain forms of intelligent reasoning (e.g. debugging failed reasoning processes) would be impossible.

The new kinds of machine table corresponding to a TM's machine table would have to be able to manipulate information about continuously varying structures, e.g. comparing two such processes and discovering how their consequences differ. Later meta-meta-...cognitive mechanisms, might add a collection of new forms of intelligence.

What replaces the Turing Machine table?

A Turing machine has a central mechanism that can be specified by a finite state transition graph with labelled links specifying actions to be performed on the tape. The proposed Super-Turing machine will need something a lot more complex, not restricted to discrete states, allowing some intrinsic (i.e. not simulated) parallelism so that some processes of change can observe other processes as they occur.

The central machine inspecting and changing the Super-Turing membrane, cannot be discrete insofar as it has to encode recognition of both non-discrete static differences (e.g. a narrowing gap between two lines as measured increasingly close to a intersection point), and recognition of continuous changes in time, including in some cases comparisons of rates of change.

For example, the deforming triangle example in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html
involves consideration of what happens as a vertex of a triangle moves along a straight line, steadily increasing its distance from the opposite side of the triangle. The machine needs to be able to detect steady increase or decrease in length or size of angle, but does not require use of numerical measures or probabilities, or changing probabilities.

So computation of derivatives, i.e. numerical rates of change, etc. need not be relevant except in very special cases, including cases of online control of actions, discussed by Gibson. Many examples of mathematical discovery seem to arise from offline uses of these mechanisms, to reason about possible actions and consequent changes without actually performing the actions. And many such cases will make use categorisations and partial orderings of structures and processes rather than measures.

So the machine's "brain" mechanisms, need to be able to make use of categorisations like 'static', 'increasing', 'decreasing', 'a is closer to b than to c' etc., including making use of partial orderings of structures and processes (i.e. spatial and temporal orderings). Illustrations of these ideas can be found in
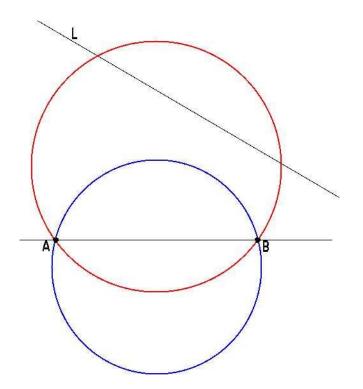http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html (or pdf)
There is also some discussion, illustrated by videos, of the kinds of intermediate information structures required for perception of changing configurations of visible parts of 3D scenes as objects move, or the perceiver moves or both move, here:
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/chairs

The mechanisms that evolved to support such perceptual processes in our ancestors and in other intelligent species (squirrels, elephants, apes, etc.) do far more than Gibson acknowledged in his theory of affordance perception. In particular I think (following Kant) that they provide the foundation for the ancient mechanisms of mathematical discovery that made possible the achievements of Archimedes, Euclid, Zeno, etc.

Possible changes to a spatial configuration: The Apollonius problem

Figure: Points Lines and Circles

Consider the two points indicated by A and B and the line L. Two circles are shown that pass through points A and B. It is clear that there are many more such circles -- infinitely many of them if space is infinitely divisible. Some circles through A and B do not intersect L, e.g. the blue circle. Other circles through A and B intersect L in two places, e.g. the red circle. Is there a circle passing through points A and B that touches L at only one point, so that L is a tangent to the circle? How do you know?

What kind of brain machinery makes it possible to reason that there must be such a circle? This is a problem from ancient geometry, and there is a construction for finding the circle that passes through A and B and meets L as a tangent, known as Apollonius' construction (also mentioned in the next section).
See http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html

As shown there, if the line L is not parallel to AB there will be two circles passing through points A and B that meet L as a tangent. One of the circles has a center above the line AB and one has a centre below.

I shall later add some thoughts here about how the need to be able to perform these spatial reasoning tasks suggests requirements for our Super-Turing machinery.

Impossible routes/trajectories

It is not very difficult to write a program that solves mazes, e.g. by always turning left at a junction (provided the maze is fully connected). However if the maze terrain is implemented as a digitized very high resolution image then such a maze-searching program might miss very small gaps in walls, or very narrow channels if its motion uses discrete steps that are larger than the smallest gaps.

When a maze program actually finds a continuous route between its start location and a specified target location then that shows that the maze has at least one solution. But if it fails to find a continuous route that may simply be due to limitations of the search strategy used, including the possibility that it has missed a very narrow gap.

This illustrates a general point: proving that a certain type of entity is possible is, in many cases, much easier than discovering impossibility, i.e. that there cannot be any instances of that type. That's because finding any instance of the type proves possibility, whereas proving impossibility (or necessity) requires more powerful cognitive resources: i.e. some way of exhaustively specifying locations in a possibility space so that they can all be shown to satisfy or not to satisfy some condition.

This is sometimes easy for simple, discrete, possibility spaces, e.g. the space of combinations of truth-values for an expression in propositional calculus with a fixed set of boolean variables. Although the number of combinations expands exponentially with the number of variables, it is always finite, whereas the set of continuous paths between two points in a 2D or 3D space is typically infinite, unless some special constraint is specified (e.g. being straight, or being a circle with centre at a third specified point).
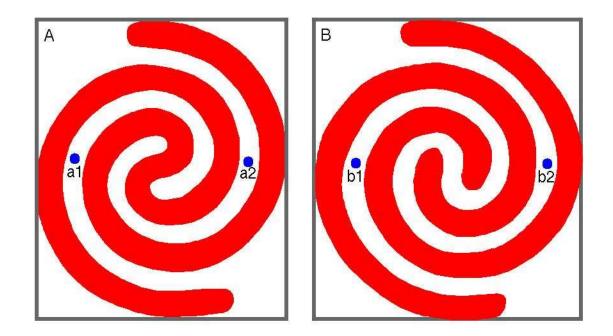
Note that although finding an instance conclusively proves possibility, there are branches of mathematics, engineering and science where finding an instance may be very difficult. E.g. although every mathematical proof using standard logic, algebra and arithmetic is finite, the space of possible proofs is unbounded, so finding a proof that actually exists may require a very long search. If there is no such proof the search will continue forever.

This is also true in geometry, since some simply described spatial configurations may require complex constructions, e.g. the statement above that, given two points A and B and a line L distinct from the line AB, find a circle C such that C passes through points A and B, and has L as a tangent. For details see
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html

A different example involves answering a question about these two figures that can be thought of as requiring consideration of an infinite collection of possibilities, without going through infinitely many steps:

Figure: Spirals

Is there a continuous route through the white space in **A** joining blue dots at **a1** and **a2**?
Is there a continuous route through the white space in **B** joining blue dots at **b1** and **b2**?

Consider the questions about existence of possible routes in Fig: Spirals A, and B. In both cases the route should not include any of the red space. The answer "yes" is easy to defend if a route has been found.

If the routes are thought of as arbitrarily thin paths linking the two points, the ability to detect when the answer is "No" is much harder to explain, as it requires an ability to survey completely a potentially infinite space of possible routes, What sort of brain mechanism, or simulated brain mechanism can provide that ability, or an equivalent ability avoiding explicit consideration of an infinite set of possibilities?

All of this is crucial to some of the uses of visual sensing (or other spatial sensing) in more or less fine-grained online control (emphasised by James Gibson), as opposed to the use of vision to categorise, predict, explain, etc. Some examples involving affordance detection going beyond Gibsonian online control are discussed in Sloman (2007-14).

For now I wish to focus mainly on the role of impossibility detection in mathematical reasoning. The difference between existence and non-existence of a route linking two blue dots without ever entering a red area is a mathematical difference. I expect most readers will not have much difficulty deciding whether such a route exists in Figure A or Figure B.

What sort of brain mechanisms can perform an exhaustive search of all possible routes from A that do not enter any red space and discover no such route reaches location B in one of the pictures? Does it really involve checking infinitely many possible routes starting fom one of the blue dots?

How can brain mechanisms implement such an exhaustive checking process, covering an enormous, possibly infinite, variety of cases. For now I'll leave that question for readers to think about.

Further requirements for intelligent, mathematical perception

A deep, and difficult, requirement for the proposed machine is that it needs to be able to detect that direction of change of one feature, e.g. increasing height of a fixed base triangle, seems to be correlated with direction of change of another, e.g. decreasing size of the angle. This requires use of partial ordering relations (getting bigger, getting smaller) and does not require use of numerical measurements.

It is also not a statistical correlation found in collections of data by statistical analysis. It's a perceived feature of a process in which two things necessarily change together.

Moreover mathematical consciousness involves seeing why such relationships *must* hold.

> I have a large, and growing, collection of examples. Many more, related to perception of possibilities and impossibilities, are collected in
> http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html
> and in documents linked from there.

It is not too hard to think of mechanisms that can observe such correspondences in perceived processes, e.g. using techniques from current AI vision systems. These are relatively minor extensions of mechanisms that can compare length, area, orientation, or shape differences in static structures without using numerical measurements.

What is much harder is explaining how such a Super-Turing mechanism can detect a necessary connection between two structures or processes.

The machine needs to be able to build bridges between the two detected processes that "reveal" an invariant structural relationship.

In Euclidean geometry studied and taught by human mathematicians, construction-lines often build such bridges, for example in standard proofs of the triangle sum theorem, or Pythagoras' theorem. (A taxonomy of cases is needed.)

But, it is important to stress that these mechanisms are not infallible. For example this document explains how I was misled at first by the stretched triangle example, because I did not consider enough possible lines of motion for the triangle's vertex.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/deform-triangle.html

The proposed machine will almost certainly generate the sorts of mistakes that Lakatos documented (in his *Proofs and Refutations*), and much simpler mistakes that can occur in everyday mathematical reasoning).

But it must also have the ability to detect and correct such errors in at least some cases -- perhaps sometimes with help of social interactions that help to expand the ideas under consideration, e.g. by merging two or more lines of thought, or sets of examples explored by different individuals.

I expect some computer scientists/AI theorists would not be happy with such imperfections: they would want a mathematical reasoning/discovery machine to be infallible.

But that's clearly not necessary for modelling/replicating human mathematical minds. Even the greatest mathematicians can make mistakes: they are not infallible.

(Incidentally this dispenses with much philosophical effort on attempting to account for infallibility, e.g. via 'self-evidence'.)

All this needs to be put into a variety of larger (scaffolding) contexts:

- evolution of the basic mechanisms (e.g. because they are relevant to intelligent control of actions)
- evolution of the meta-cognitive mechanisms that can *notice* patterns of discovery, error, error correction, debugging, plan repair...
- evolution of communication abilities based on helping others control or use their metacognition ... etc.
- evolution of external aids (extended genotype/phenotype???)
- (social/cultural) evolution of practices of education, organisation of knowledge, etc. --> Euclid's Elements

But there are many unknown details, including which (human and non-human) brain mechanisms can do such things and (a) how and when they evolved and (b) how and when they develop within individuals.

Properties of the Super Turing membrane machine
21 Nov 2017: moved here from another document.

The conjectured virtual membrane will have various ways of acquiring "painted" structures, including the following (a partial, provisional, list):

- static projections from what Gibson called the "optic array" -- e.g. when you see the outline of an object;
- dynamic projections from perceived paths of moving objects (leaving a subsequently accessible "visual trail" as sometimes depicted in cartoon drawings of moving objects);
- ways of grouping and ungrouping items perceived as moving together or separately;
- ways of grouping and ungrouping items intended to be moved together or separately;
- ways of deforming shapes (e.g. bend a line, move one vertex of a triangle, or make a dent or bulge in the side of a circle);
- ways of checking relationships such as contact, containment relative size (of lengths, areas, distance from perceiver);
- ways of managing two or more such surfaces concurrently, e.g. using one as a fixed background against which another can move, either in the same plane or by moving towards or away from the background;
- ways of abstracting from details of the contents painted, thereby producing new idealisations, e.g. thinking about perfectly thin lines, perfectly straight lines, perfectly symmetrical shapes, etc. although it is impossible to identify them visually (e.g. distinguish them from very closely

similar lines) or to create them using line drawing tools or even to imagine them with the precision implied by their definition.

NB: regarding the last point: I am not suggesting that evolution actually produced perfectly thin, perfectly straight, structures, or mechanisms that could create such things. Rather the mechanisms would have the ability to (implicitly or explicitly) *postulate* such limiting case features, *represent* them, and *reason about* their relationships, by extrapolating from the non-limiting cases. (Is this what Kant was saying in [1781](#)?)

For example, the assumption that when two perfectly thin lines cross the intersection is a point with zero diameter would depend on an ability to extrapolate from what happens if two lines with non-zero thickness cross, and then they both become narrower and narrower, so that the overlap location becomes smaller in all directions. (Yet another "theorem" lurks there!)

It should not be assumed that any of this produces high precision projections, though a conjectured learning process (possibly produced by evolution across many stages of increasing sophistication) may generate mechanisms that can "invent" limiting cases, e.g. perfectly thin lines, perfectly straight lines, etc. perhaps starting with simpler versions of the constructions presented by Dana Scott in [Scott(2014)](#).

Monitoring mechanisms

In addition, a repertoire of "meta-membrane" monitoring mechanisms needs to be available that can detect potentially useful (or even merely interesting!) changes, including items coming into contact or moving apart, orderings being changed, new sub-structures created or disassembled, detection of new relations arising, or old relations being destroyed (contact, ordering, etc.), when membrane manipulation processes occur.

I have not attempted to answer the question whether the proposed membrane mechanisms (still under-specified) require new kinds of information processing (i.e. computation in a general sense) that use physical brain mechanisms that are not implementable on digital computers, perhaps because they rely on a kind of mixture of continuity and discreteness found in many chemical processes in living organisms.

It could turn out that everything required is implementable in a suitable virtual machine implemented on a digital computer. For example, humans looking at a digital display may perceive the lines, image boundaries and motions as *continuous* even though they are in fact *discrete*. This can happen when a clearly digital moving display is viewed through out of focus lenses, or at a distance, or in dim light, etc. In that case the blurring or smoothing is produced by physical mechanisms before photons hit the retina.

But it is also possible to treat a digital display as if it were discrete, for example assigning sub-pixel coordinates to parts of visible lines, or motion trajectories. That sort of blurring loses information, but may sometimes make information more useful, or more tractable. It could be useful to build visual systems for robots with the ability to implement various kinds of virtual de-focusing mechanisms for internal use when reasoning about perceived structures, or controlling actions on perceived structures, e.g. moving a hand to pick up a brick.

Insofar as human retinas have concentric rings of feature detectors with higher resolution detectors near a central location (the fovea) and lower resolution detectors further from the centre, it can be viewed as a mechanism that gives perceivers the ability to change the precision at which certain features in the optic array are focused. It may have other applications.

Sub-neural chemical computations?

In theory, no new kinds of physical machine would be needed if the membrane mechanisms can use new kinds of digitally implementable virtual machinery, using virtual membranes and membrane operations. However, even if that is theoretically possible it may be intractable if the numbers of new virtual machine components need to match not neural but sub-neural molecular/chemical computational resources in brains. The number of transistors required to model such a mechanism digitally might not fit on our planet.

Compare the challenges to conventional thinking about brains implicit in Schrödinger(1944) and explicit in Grant(2010), Gallistel&Matzel, 2012 and Trettenbrein(2016), suggesting that important aspects of natural information processing are chemical, i.e. sub-neural.

It is possible that such molecular-level forms of information processing could be important for the sorts of information processing brain functions postulated in Trehub(1991), though molecular level implementation would require significant changes to Trehub's proposed implementation of his ideas.

Turing's 1952 paper on chemistry-based morphogenesis Turing (1952) at first sight appears to be totally unconnected with his work on computation (except that he mentioned using computers to simulate some of the morphogenesis processes). But if Turing had been thinking about requirements for replicating geometrical and topological reasoning used by ancient mathematicians and learners today, then perhaps he thought, or hoped, the chemical morphogenesis ideas would be directly relevant to important forms of computation, in the general sense of information processing. In that case his ideas might link up with the ideas about sub-neural computation referenced above, which might perhaps play a role in reasoning mechanisms conjectured in Sloman(2007-14), which draws attention to kinds of perception of topology/geometry based possibilities and impossibilities that were not, as far as I know, included in the kinds of affordance that Gibson considered.

The membrane (or multi-membrane) machine needs several, perhaps a very large number, of writeable-readable-sketchable surfaces that can be used for various purposes, including perceiving motion, controlling actions, and especially considering new possibilities and impossibilities (proto-affordances). Some examples of transitions that need to be coped with are in the short video scenarios here:
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/chairs

The idea also needs to be generalised to accommodate inspectable 3D structures and processes, like nuts rotating on bolts, as discussed another document. (Something about this may be added here later.)

The brain mechanisms to be explained are also likely to have been used by ancient mathematicians who made the amazing discoveries leading up to publication of Euclid's *Elements*, and later.   http://www.gutenberg.org/ebooks/21076

I think there are deep connections between the abilities that made those ancient mathematical discoveries possible, and processes of perception, action control, and reasoning in many intelligent organisms, as suggested in the workshop web page mentioned above.

One consequence of the proposal is that Euclid's axioms, postulates and constructions are not arbitrarily adopted logical formulae in a system that implicitly defines the domain of Euclidean geometry. Neither are they mere empirical/statistical generalisations capable of being refuted by new observations.

Rather, as Kant suggested, they were all mathematical discoveries, using still unknown mechanisms in animal brains originally produced (discovered) by evolution with functions related to reasoning about perceived or imagined spatial structures and processes, in a space supporting smoothly varying sets of possibilities, including continuously changing fields of view, visible portions of surfaces, shapes, sizes, orientations, curvature and relationships between structures, especially partial orderings (e.g. comparisons of size, containment, angle, curvature, etc.)

Despite all the smooth changes, the space also supports many interesting emergent discontinuities and invariants, that the ancients discovered and discussed, many of which seem to be used unwittingly(?) by other intelligent species and pre-verbal children, for example, a smoothly changing line can change its curvature at a certain location from being in one direction to being in the opposite direction (e.g. from left to right). The same thing can happen if focus of attention is moved along a fixed line from a location where the line curves to the left to a location where it curves in the opposite direction. Fast moving animals need to be able to detect such changes in paths or trajectories they are following, altering the forces they apply to achieve the desired change of direction -- e.g. a bird swerving between branches in a tree to get to its nest.

Adult humans normally have additional meta-cognitive mechanisms, though it is not clear what they all are, nor how and when they develop.

A few examples used in much every-day visual perception, action control and planning are discussed in Sloman(2007-14)).

The location at which the size of a vertex of a triangle peaks when moved along a straight line that does not pass through the base of the triangle is another, surprisingly complicated, example, discussed in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/apollonius.html

---

Is chemistry essential for some animal competences?

I suspect chemistry-based reasoning mechanisms are important in brains, as Turing suggested in his Mind 1950 paper, though the comment usually goes unnoticed. I think Kenneth Craik had related (under-developed) ideas in 1943, e.g. wondering how a tangle of neurons could represent straightness or triangularity... before digital computers and digitised image arrays had been invented.

I conjecture that Turing may have been thinking about these issues when he wrote his paper on morphogenesis, published two years before he died: Turing (1952). For a useful summary for non-mathematicians, see Ball (2015)

TO BE CONTINUED

This needs to be related to the theory of evolved construction kits:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html
    Construction kits for evolving life
    (Including evolving minds and mathematical abilities.)
An older version of the construction kits paper (frozen mid 2016) was published in Cooper and Soskova (Eds) 2017.

REFERENCES AND LINKS

- Philip Ball, 2015, "Forging patterns and making waves from biology to geology: a commentary on Turing (1952) 'The chemical basis of morphogenesis'",
  http://dx.doi.org/10.1098/rstb.2014.0218

- Dana Ballard and Chris Brown, 1982, *Computer Vision,* Prentice Hall, Englewood Cliffs, New Jersey 07632, USA, online at:
  http://homepages.inf.ed.ac.uk/rbf/BOOKS/BANDB/bandb.htm

- John Seely Brown, Richard R. Burton, and Kathy M. Larkin, 1977,
  Representing and using procedural bugs for educational purposes, *Proceedings of the 1977 annual conference,* ACM '77, pp. 247--255, ACM, New York, NY, USA,
  http://doi.acm.org/10.1145/800179.810211

- S. Barry Cooper and Mariya I. Soskova, editors 2017 The Incomputable: Journeys Beyond the Turing Barrier Springer-Verlag,
  http://www.springer.com/gb/book/9783319436678

- Richard Dawkins, 1986, *The Blind Watchmaker: Why the Evidence of Evolution Reveals a Universe without Design*
  Norton & Company, Inc,

- Brian V. Funt (1980), Problem-Solving with Diagrammatic Representations, in *Artificial Intelligence* 13 (1980), pp201-230 North-Holland
  http://www.cs.sfu.ca/~funt/ProblemSolvingWithDiagrammaticRepresentations_AIJournal1980.pdf

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

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

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

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

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

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

- Much of Jean Piaget's work is also relevant, especially his last two (closely related) books written with his collaborators:
  Possibility and Necessity
  Vol 1. The role of possibility in cognitive development (1981)
  Vol 2. The role of necessity in cognitive development (1983)
  Tr. by Helga Feider from French in 1987

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

- Aaron Sloman (1962) *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth* DPhil thesis, Oxford University, May 1962. (This was transcribed and made available as searchable text in 2016.. I am very grateful to Luc Beaudoin for locating the transcription service and for much help with checking and correcting the transcription.)
  http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962

- A. Sloman, 1984 (extended later) Experiencing Computation: A Tribute to Max Clowes, originally in *New horizons in educational computing,* Ellis Horwood Series In AI Ed. Masoud Yazdani, pp. 207--219, Chichester,
  Online version with expanded obituary and biography
  http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61

- Aaron Sloman, 1971, Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence, *Proc 2nd IJCAI, London, 1971*
  Reprinted in *Artificial Intelligence,* vol 2, 3-4, pp 209-225, 1971,
  http://www.cs.bham.ac.uk/research/cogaff/62-80.html#1971-02

- Aaron Sloman, 1978 *The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind,*
  Harvester Press, 1978, http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp (Revised 2018)

- Aaron Sloman (2007-14), Predicting Affordance Changes (Alternative ways to deal with uncertainty), Unpublished discussion paper (HTML), School of Computer Science, University of Birmingham, Installed 2007, later updated
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html

- Alan Turing's 1938 thoughts on mathematical reasoning, with some comments by A. Sloman http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html (also pdf)

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

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

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

- A. M. Turing, 1939, Systems of Logic Based on Ordinals, *Proc. London Mathematical Society,* pp. 161-228, https://doi.org/10.1112/plms/s2-45.1.161
  Also reprinted in *Alan Turing: His work and impact* Elsevier 2013.

- A.M. Turing (1952) The Chemical Basis Of Morphogenesis. *Phil Trans R Soc London* B 237 237:37-72
  Note: A presentation of Turing's main ideas for non-mathematicians can be found in Ball, 2015.

- Eds. S. Barry Cooper and J. van Leeuwen, *Alan Turing: His Work and Impact*, 2013, Elsevier, Amsterdam,

  A.A.S Weir, J. Chappell, and A. Kacelnik,(2002), Shaping of hooks in New Caledonian crows, *Science,* vol 297, page 981.
  (The videos on the laboratory web site show more complex and varied solutions than the paper reports.)

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