

Natural Vision and Mathematics: Seeing Impossibilities

Bridging the Gap Workshop IJCAI 2016(Furbach & Schon, 2016)

Aaron Sloman*

School of Computer Science, University of Birmingham, UK
<http://www.cs.bham.ac.uk/~axs>

Abstract. The Turing-inspired Meta-Morphogenesis project investigates forms of biological information processing produced by evolution since the earliest life forms, especially information processing mechanisms that made possible the deep mathematical discoveries of Euclid, Archimedes, and other ancient mathematicians. Such mechanisms must enable perception of possibilities and constraints on possibilities – types of affordance perception not explicitly discussed by Gibson, but suggested by his ideas. Current AI vision and reasoning systems lack such abilities. A future AI project might produce a design for “baby” robots that can “grow up” to become mathematicians able to replicate (and extend) some of the ancient discoveries, e.g. in the way that Archimedes extended Euclidean geometry to make trisection of an arbitrary angle possible. This is relevant to many kinds of intelligent organism or machine perceiving and interacting with structures and processes in the environment. This would demonstrate the need to extend Dennett’s taxonomy of types of mind to include Euclidean (or Archimedean) minds, and would support Immanuel Kant’s philosophy of mathematics.

Keywords: AI, Kant, Mathematics, Meta-morphogenesis, intuition, Euclid, Geometry, Topology, Kinds-of-minds, Meta-cognition, Meta-meta-cognition, etc.

Mathematics and computers

It is widely believed that computers will always outperform humans in mathematical reasoning. That, however, is based on a narrow conception of mathematics that ignores some of the history of mathematics, e.g. achievements of Euclid and Archimedes, and also ignores kinds of mathematical competence that are a part of our everyday life, but mostly go unnoticed, e.g. topological reasoning abilities.¹ These are major challenges for AI, especially AI projects to replicate or model human mathematical competences. I don’t think we are ready to build working systems with these competences, but I’ll introduce a research programme that may, eventually, lead us towards adequate models. It challenges common assumptions about the nature and evolution of human language and hints at evolved mechanisms supporting mathematical competences and closely related perceptual and reasoning competences.

Formal mechanisms based on use of arithmetic, algebra, and logic, dominate AI models of mathematical reasoning, but the great ancient mathematicians did not use modern logic and formal systems. Such things are therefore not *necessary* for mathematics. they are *part* of mathematics: a recent addition, but appear not

* This is a snapshot of part of the Turing-inspired Meta-Morphogenesis project.
<http://goo.gl/eFnJb1> (This preprint will later have more links included.)

¹ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/shirt.html>

to be needed for older forms of human and animal reasoning (e.g. topological reasoning). By studying achievements of ancient mathematicians, pre-verbal human toddlers, and intelligent non-human animals, especially perception and reasoning abilities that are not matched by current AI systems, or explained by current theories of how brains work, we can identify major challenges. Powerful new *internal* languages are needed, similar to evolved languages used for perceiving, thinking about, and reasoning about shapes, structures and spatial processes. If such internal languages are used by intelligent non-human animals and pre-verbal toddlers, their evolution must have preceded evolution of human languages for communication (Sloman, 1978b, 1979, 2015). Structured *internal* languages (for storing and using information) must have evolved *before* external languages, since there would be nothing to communicate and no use for anything communicated, without pre-existing internal mechanisms for constructing, manipulating and using structured percepts, intentions, beliefs, plans, memorised maps, etc.

The simplest organisms (viruses?) may have only passive physical/chemical reactions, and only trivial decisions and uses of information (apart from genetic information). Slightly more complex organisms may use information only for taking Yes/No or More/Less or Start/Stop decisions, or perhaps selections from a pre-stored collection of possible internal or external actions. (Evolution's menus!) More complex internal meaning structures are required for cognitive functions based on information contents that can vary in structure and complexity, like the Portia spider's ability to study a scene for about 20 minutes and then climb a branching structure to reach a position above its prey, and then drop down for its meal (Tarsitano, 2006). This requires an initial process of information collection and storage in a scene-specific structured form that later allows a pre-computed branching path to be followed even though the prey is not always visible during the process, and portions of the scene that are visible keep changing as the spider moves. Portia is clearly conscious of much of the environment, during and after plan-construction. As far as I know, nobody understands in detail what the information processing mechanisms are that enable the spider to take in scene structures and construct a usable 3-D route plan, though we can analyse the computational requirements on the basis of half a century of AI experience.

This is one example among many cognitive functions enabling individual organisms to deal with static structured situations and passively perceived or actively controlled processes, of varying complexity, including control processes in which parts of the perceiver change their relationships to one another (e.g. jaws, claws, legs, etc.) and relationships to other things in the environment (e.g. food, structures climbed over, or places to shelter).

Abilities to perceive plants in natural environments, such as woodlands or meadows, and, either immediately or later, make use of them, also requires acquisition, storage and use of information about complex objects of varying structures, and information about complex processes in which object-parts change their relationships, and change their visual projections as the perceiver moves. Acting on perceived structures, e.g. biting or swallowing them, or carrying them to a part-built nest to be inserted, will normally have to be done differently in

different contexts, e.g. adding twigs with different sizes and shapes at different stages in building a nest. How can we make a robot that does this?

Conjecture:

Non-human abilities to create and use information structures of varying complexity are evolutionary precursors of human abilities to use grammars and semantic rules for languages in which novel sentences are understood in systematic ways to express different more or less complex percepts, intentions, or plans to solve practical problems, e.g. using a lexicon, syntactic forms, and compositional semantics. In particular, a complex new information structure can be assembled and stored, then later used as an information structure (e.g. plan, hypothesis) while selecting or controlling actions.

We must not, of course, be deceived by organisms that *appear* to be intentionally creating intended structures but are actually doing something much simpler that creates the structures as a by-product, like bees huddled together, oozing wax, vibrating, and thereby creating a hexagonal array of cavities, that look designed but were not. Bees have no need to count to six to do that. Most nest-building actions, however, are neither random nor fixed repetitive movements. They are guided, in part, by missing portions of incomplete structures, where what's missing and what's added keeps changing. So the builders need internal languages with generative syntax, structural variability, (context sensitive) compositional semantics, and inference mechanisms in order to be able to encode all the relevant varieties of information needed. Nest building competences in corvids and weaver birds are examples. Human architects are more complicated!

Abilities to create, perceive, change, manipulate, or use meaning structures (of varying complexity) enable a perceiver of a novel situation to take in its structure and reason hypothetically about effects of possible actions - without having to collect evidence and derive probabilities. The reasoning can be geometric or topological, without using any statistical evidence: merely the specification of spatial structures. Reasoning about what is *impossible* (not merely improbable) can avoid wasted effort. For example, the "polyflap" domain² was proposed as an artificial environment illustrating some challenging cognitive requirements. It contains arbitrary 2D polygonal shapes each with a single (non-flat) fold forming a new 3D object. An intelligent agent exploring polyflaps could learn that any object resting on surfaces where it has a total of two contact points can rotate in either direction about the line joining the contact points. Noticing this should allow the agent to work out that achieving stability would require a supporting surface on which at least one more part of the supported object can rest. In simple cases all three points are on one horizontal plane: e.g. on a floor. But an intelligent agent that understands stability should be able to produce stability with three support points on different, non-co-planar surfaces, e.g. the tops of three pillars with different heights. Any two of the support points on their own would allow tilting about the line joining the points. But if the third support point is not on that line, and a vertical line through the object's centre of gravity

² <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps> (Sloman, 2005)

goes through the interior of the triangle formed by the support points then the structure will be stable³.

An intelligent machine should be able to reason in similar ways about *novel* configurations. This illustrates a type of perception of affordances in the *spirit* of Gibson's theory, but going beyond examples he considered. I don't recall him mentioning use of geometrical or topological reasoning in deciding what would be stable. This contradicts a common view that affordances are discovered through statistical learning. Non-statistical forms of reasoning about affordances in the environment (possibilities for change and constraints on change) may have been a major source of the amazing collection of discoveries about topology and geometry recorded in Euclid's *Elements*. Such forms of discovery are very important, but still unexplained.

Intelligent non-human animals (e.g. Portia), and humans, use evolved mechanisms that can build, manipulate and use *structured* internal information records whose required complexity can vary and whose information *content* is derivable from information about parts, using some form of "compositional semantics", as is required in human spoken languages, logical languages, and programming languages. However, the internal languages need not use *linear* structures, like sentences. In principle they could be trees, graphs, nets, map-like structures or types of structure we have not yet thought of. Human sign languages are not restricted to discrete temporal sequences of simple signs: usually movements of hands, head and parts of the face go on in parallel. This may be related to use of non-linear internal languages for encoding perceptual information, including changing visual information about complex structured scenes and tactile information gained by manual exploration of structured objects. In general the 3-D world of an active 3-D organism is not all usefully linearizable.

The variety of types of animal that can perceive and act intelligently in relation to novel perceived environmental structures, suggests that many use "internal languages" in a generalised sense of "language" ("Generalised Languages" or GLs), with *structural variability* and *context sensitive compositional semantics*, which must have evolved long before human languages were used for communication (Sloman & Chappell, 2007)⁴. The use of external, structured, languages for communication presupposes internal perceptual mechanisms using (GLs), e.g. for parsing messages and relating them to percepts and intentions. There are similar requirements for intelligent nest building by birds and for many forms of complex learning and problem solving by other animals, including elephants, squirrels, cetaceans, monkeys, apes and pre-verbal human toddlers⁵.

Is there a circularity?

In the past, philosophers would have argued (scornfully!) that postulating the need for an internal language IL to be used in understanding an external language EL, would require yet another internal language for understanding IL, and so on, leading to an infinite regress. But AI and computer systems engineering

³ I first discovered this "Polyflap stability theorem" while writing this paper.

⁴ See <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>

⁵ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>

demonstrate that there need not be an infinite regress. Our understanding has expanded enormously since the 1960s as AI and software engineering researchers should know. How brains achieve similar feats is unknown, however.

Mechanisms for communicating information using familiar human languages would have no point if senders and recipients were not already information *users*, otherwise they would have nothing to communicate, and would have no way to *use* information communicated to them. Yet there is much resistance to the idea that rich internal languages used for *non-communicative* purposes evolved before communicative languages. That may be partly because many people do not understand the computational requirements for the competences displayed by pre-verbal humans and other animals, and partly because they don't understand complex virtual machines do not require an infinite hierarchy of internal languages.

These comments about animals able to perceive, manipulate and reason about varied objects and constructions, apply also to pre-verbal human toddlers playing with toys and solving problems, including manipulating food, clothing, and even their parents. (See also Note 5.)

The full repertoire of biological vehicles and mechanisms for information bearers must include both mechanisms and meta-mechanisms (mechanisms that construct new mechanisms) produced by natural selection and inherited via genomes, and also individually discovered/created mechanisms, in humans, and other altricial species with “meta-configured” competences in the terminology of (Chappell & Sloman, 2007).

Dennett (e.g. 1995) firmly resists the internal language idea: he argues, on the contrary, that consciousness followed evolution of language mechanisms used for external communication, later also used internally for silent self-communication. But Portia spiders surely did not need ancestors that discussed planned routes to capture prey before they evolved the ability to reason silently about terrain and routes and plans.

Creation vs Learning:

Evidence from deaf children in Nicaragua (Senghas, 2005) and subtle clues in non-deaf children, show that children do not *learn* languages from existing users. Rather, they develop mechanisms of increasing power, that enable them to *create* languages collaboratively⁶. Normally children collaborate as a relatively powerless minority, so the creation process looks like imitative learning.

Many familiar processes of perceiving, learning, intending, reasoning, planning, plan execution, plan debugging, etc. would be impossible if humans (and perhaps some other intelligent animals) did not have rich *internal* languages and language manipulation mechanisms (GL competences). For more on this see Note 4. Some philosophers, cognitive scientists and others believe that structure-based “old fashioned” AI has failed. But the truth is that NO form of AI has “succeeded”, apart from narrowly focused AI applications. Many of them

⁶ This video gives some details: <https://www.youtube.com/watch?v=pjtioIFuNf8>

are very shallow and more restricted than symbolic planners and reasoners⁷. Compare (Rescorla, 2015).

We still need to learn much more about the nature of internal GLs, the mechanisms required, and their functions in various kinds of intelligent animal. We should not expect them to be much like kinds of human languages or computer languages we are already familiar with, if various GLs also provide the internal information media for perceptual contents of intelligent and fast moving animals like crows, squirrels, hunting mammals, monkeys, and cetaceans. Taking in information about rapidly changing scenes, needs something different from Portia's internal language for describing a fixed route. Moreover, languages for encoding information about changing visual contents will need different sorts of expressive powers from languages for human conversation about the weather or the next meal⁸.

Many people have studied and written about various aspects of non-verbal communication and reasoning, including, for example, contributors to (Glasgow, Narayanan, & Chandrasekaran, 1995), and others who have presented papers on diagrammatic reasoning, or studied the uses of diagrams by young children. But there are still deep gaps, especially related to mathematical discoveries.

Many of Piaget's books provide examples, some discussed below. He understood better than most what needed to be explained, but he lacked any understanding of programming or AI and he therefore sought explanatory models where they could not be found, e.g. in boolean algebras and group theory.

The importance of Euclid for AI

AI sceptics attack *achievements* of AI, whereas I am attacking the *goals* of researchers who have not noticed the need to explain some very deep, well known but very poorly understood, human abilities: the abilities that enabled our ancestors prior to Euclid, without the help of mathematics teachers, to make the sorts of discoveries that eventually stimulated Euclid, Archimedes and other ancient mathematicians who made profound non-empirical discoveries, leading up to what is arguably the single most important book ever written on this planet: Euclid's *Elements*⁹. Thousands of people all around the world are still putting its discoveries to good use every day even if they have never read it¹⁰.

⁷ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/chewing-test.html>

⁸ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision/plants> presents a botanical challenge for vision researchers.

⁹ There seems to be uncertainty about dates and who contributed what. I'll treat Euclid as a figurehead for a tradition that includes many others, especially Thales, Pythagoras and Archimedes - perhaps the greatest of them all, and a mathematical precursor of Leibniz and Newton. (More names are listed here: https://en.wikipedia.org/wiki/Chronology_of_ancient_Greek_mathematicians.) I don't know much about mathematicians on other continents at that time or earlier. I'll take Euclid to stand for all of them, because of the book that bears his name.

¹⁰ Moreover, it does not propagate misleading falsehoods, condone oppression of women or non-believers, or promote dreadful mind-binding in children.

As a mathematics graduate student interacting with philosophers around 1958, I formed the impression that the philosopher whose claims about mathematics were closest to what I knew about *doing* mathematics, especially geometry, was Immanuel Kant. But his claims about our knowledge of Euclidean geometry seemed to have been contradicted by recent theories of Einstein and empirical observations by Eddington. Philosophers therefore thought that Kant had been refuted, ignoring the fact that Euclidean geometry without the parallel axiom remains a deep and powerful body of geometrical and topological knowledge, and provides a basis for constructing three different types of geometry: Euclidean, elliptical and hyperbolic, the last two based on alternatives to the parallel axiom¹¹. We'll see also that it also has an extension that makes trisection of an arbitrary angle possible, unlike pure Euclidean geometry. These are real mathematical discoveries about a type of space, not about logic, and not about observed statistical regularities.

My experience of doing mathematics suggested that Kant was basically right in his claims against David Hume: many mathematical discoveries provide knowledge that is *non-analytic* (i.e. synthetic, not proved solely on the basis of logic and definitions), *non-empirical* (i.e. possibly triggered by experiences, but not based on experiences, nor subject to refutation by experiment or observation, if properly proved), and *necessarily true* (i.e. incapable of having counter-examples, not contingent).

This does not imply that human mathematical reasoning is infallible: (Lakatos, 1976) demonstrated that even great mathematicians can make various kinds of mistakes in exploring something new and important. Once discovered, mistakes sometimes lead to new knowledge. So a Kantian philosopher of mathematics need not claim that mathematicians produce only valid reasoning¹².

Purely philosophical debates on these issues can be hard to resolve. So when Max Clowes¹³ introduced me to AI and programming around 1969 I formed the intention of showing how a baby robot could grow up to be a mathematician in a manner consistent with Kant's claims. But that has not yet been achieved. What sorts of discovery mechanisms would such a robot need?

Around that time, a famous paper by McCarthy and Hayes claimed that **logic** would suffice as a form of representation (and therefore also reasoning) for intelligent robots. The paper discussed the representational requirements for intelligent machines, and concluded that “... *one representation plays a dominant role and in simpler systems may be the only representation present. This is a representation by sets of sentences in a suitable formal logical language... with function symbols, description operator, conditional expressions, sets, etc.*” They discussed several kinds of adequacy of forms of representation, including metaphysical, epistemological and heuristic adequacy (vaguely echoing distinctions Chomsky had made earlier regarding types of adequacy of linguistic theories).

¹¹ <http://web.mnstate.edu/peil/geometry/C2EuclidNonEuclid/8euclidnoneuclid.htm>

¹² My 1962 DPhil thesis presented Kant's ideas, before I had heard about AI. (**Sloman, 1962**)

¹³ <http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html>

Despite many changes of detail, a great deal of important AI research has since been based on the use of logic as a GL, now often enhanced with statistical mechanisms.

Nevertheless thinking about mathematical discoveries in geometry and topology and many aspects of everyday intelligence suggested that McCarthy and Hayes were wrong about the sufficiency of logic. I tried to show why in (Sloman, 1971) and later papers. Their discussion was more sophisticated than I have indicated here. In particular, they identified different sorts of criteria for evaluating forms of representation, used for thinking or communicating:

A representation is metaphysically adequate if the world could have that form without contradicting relevant facts about aspects of reality.

A representation is called epistemologically adequate for a person or machine if it can be used practically to express the facts that one actually has about the aspect of the world.

A representation is called heuristically adequate if the reasoning processes actually gone through in solving a problem are expressible in the language.

Ordinary language is obviously adequate to express the facts that people communicate to each other in ordinary language. It is not adequate to express what people know about how to recognize a particular face.

They concluded that a form of representation based on logic would be heuristically adequate for intelligent machines observing, reasoning about and acting in human-like environments. But this does not provide an explanation of what adequacy of reasoning is. For example, one criterion might be that the reasoning should be *incapable* of deriving false conclusions from true premisses.

At that time I was interested in understanding the nature of mathematical knowledge (as discussed in (Kant, 1781)). I thought it might be possible to test philosophical theories about mathematical reasoning by demonstrating how a “baby robot” might begin to make mathematical discoveries (in geometry and arithmetic) as Euclid and his precursors had. But I did not think logic-based forms of representation would be heuristically adequate because of the essential role played by diagrams in the work of mathematicians like Euclid and Archimedes, even if some modern mathematicians felt such diagrams should be replaced by formal proofs in axiomatic systems - apparently not noticing that that changes the investigation to a different branch of mathematics. The same can be said about Frege’s attempts to embed arithmetic in logic.

(Sloman, 1971) offered alternatives to logical forms of representation, including (among others) “analogical” representations that were not based on the kind of (Fregean) function/argument structure used by logical representations. Despite explicit disclaimers, the paper is often mis-reported as claiming that analogical representations are *isomorphic* with what they represent: which is clearly not generally true, since a 2-D picture cannot be isomorphic with the 3-D scene it represents, one of several reasons why AI vision research is so difficult.

A revised, extended, notion of validity of reasoning, was shown to include changes of pictorial structure that correspond to possible changes in the entities

or scenes depicted, but this did not explain how to *implement* a human-like diagrammatic reasoner in geometry or topology. 45 years later there still seems to be no AI system that is capable of discovering and understanding deep diagrammatic proofs of the sorts presented by Euclid, Archimedes and others. This is associated with inability to act intelligently in a complex and changing environment that poses novel problems involving spatial structures.

A more subtle challenge comes from the discovery known to Archimedes that there is a simple and natural way of *extending* Euclidean geometry by adding the *neusis* construction which makes it easy to trisect an arbitrary angle¹⁴.

As far as I know no current AI reasoning system could make such a discovery. It is definitely not connected with statistical learning: that would not provide insight into mathematical *necessity* or *impossibility*. The result was not derived from axioms: it showed that Euclid's axioms could be *extended*. Mary Pardoe, a former student, discovered a related but simpler extension to Euclid, allowing the triangle sum theorem to be proved without using the parallel axiom¹⁵.

As far as I know, nobody in AI has tried to implement abilities to discover Euclidean geometry, including topological reasoning, and the extensions mentioned here, in an AI system or robot with spatial reasoning abilities. I am still trying to understand why it is so difficult. (But not impossible, I hope.)

Not only competences of adult human mathematicians still await replication in AI. Many intelligent animals, such as squirrels, nest building birds, elephants and even octopuses have abilities to discover and use spatial manipulations of objects in their environment (or their own body parts) and apparently understand what they are doing. Betty, a New Caledonian crow, made headlines in 2002 when she was observed (in Oxford) making a hook from a straight piece of wire in order to extract a bucket of food from vertical glass tube (Weir, Chappell, & Kacelnik, 2002). The online videos demonstrate something not mentioned in the original published report, namely that Betty was able to make hooks in several different ways, all of which worked immediately without any visible signs of trial and error. She clearly understood what was possible, despite not having lived in an environment containing pieces of wire or any similar material (twigs either break if bent or tend to straighten when released). It's hard to believe that such a creature could be using logic, as recommended by McCarthy and Hayes. But what are the alternatives? Perhaps a better developed theory of GLs will provide the answer and demonstrate it in a running system.

The McCarthy and Hayes paper partly echoed Frege, who had argued in 1884 that arithmetical knowledge could be completely based on logic, But he denied that geometry could be (despite Hilbert's axiomatization of Euclidean geometry). (Whitehead & Russell, 1910–1913) had also attempted to show how the whole of arithmetic could be derived from logic. though Russell oscillated in his views about the philosophical significance of what had been demonstrated.

¹⁴ As demonstrated in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html>
I was unaware of this until I found the Wikipedia article in 2015:

https://en.wikipedia.org/wiki/Angle_trisection#With_a_marked_ruler

¹⁵ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html>

Frege was right about geometry: what Hilbert axiomatised was a combination of logic and arithmetic that demonstrated that arithmetic and algebra contained a *model* of Euclidean geometry based on arithmetical analogues of lines, circles, and operations on them, discovered by Descartes. But doing that did not imply that the *original* discoveries were arithmetical discoveries rather than discoveries about spatial structures, relationships and transformations. (Many mathematical domains have models in other domains.)

When the ancient geometers made their discoveries, they were not reasoning about relationships between logical symbols in a formal system or about numbers or equations. This implies that in order to build robots able to repeat those discoveries it will not suffice merely to give them abilities to derive logical consequences from axioms expressed in a logical notation, such as predicate calculus or the extended version discussed by McCarthy and Hayes.

Instead we'll need to understand what humans do when they think about shapes and the ways they can be constructed, extended, compared, etc. This requires more than getting machines to answer the same questions in laboratory experiments, or pass the same tests in mathematical examinations. We need to develop good theories about what human mathematicians did when they made the original discoveries, without the help of mathematics teachers, and without the kind of drill and practice now often found in mathematical classrooms. Those theories should be sufficiently rich and precise to enable us to produce working models that demonstrate the explanatory power of the theories.

As far as I know there is still nothing in AI that comes close to enabling robots to replicate the ancient discoveries in geometry and topology, nor any formalism that provides the capabilities GLs would need, in order to explain how products of evolution perceive the environment, solve problems, etc. Many researchers in AI, psychology and neuroscience, now think the core requirement is a shift from logical reasoning to statistical/probabilistic reasoning. But that cannot prove necessary truths. I suspect a deeper advance will come from extending techniques for reasoning about possibilities, impossibilities, changing topological relationships and the use of partial orderings (of distance, size, orientation, curvature, slope, containment, etc.)¹⁶.

What about arithmetic?

The arguments against any attempt to redefine geometry in terms of what follows from Hilbert's axioms can be generalised to argue against Frege's attempt to redefine arithmetic in terms of what follows from axioms and rules for logical reasoning. In both cases a previously discovered and partially explored mathematical domain was shown to be modelled using logic. But modelling is one thing; replicating another.

The arithmetical discoveries made by Euclid and others long before the discovery of modern logic were more like discoveries in geometry than like proofs in an axiomatic system using only logical inferences. However, arithmetical knowledge is not concerned only with spatial structures and processes. It involves

¹⁶ As suggested in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html>

general features of groups or sets of entities, and operations on them. For example, acquiring the concept of the number six requires having the ability to relate different groups of objects in terms of one-to-one correspondences (bijections). So the basic idea of arithmetic is that two collections of entities may or may not have a 1-1 relationship. If they do we could call them “equinumeric”. The following groups are equinumeric in that sense (treating different occurrences of the same character as different items).

[U V W X Y Z] [P P P P P P] [W Y Y G Q P]

If we focus on *types* of character rather than instances, then the numbers are different. The first box contains six distinct items, the second box only one type, and the third box five types. For now, let’s focus on instances not types.

The relation of equinumerosity has many practical uses, and one does not need to know anything about names for numbers, or even to have the concept of a number as an entity that can be referred to, added to other numbers etc. in order to make use of equinumerosity. For example, if someone goes fishing to feed a family and each fish provides a meal for one person, the fisherman could take the whole family, and as each fish is caught give it to an empty-handed member of the family, until everyone has a fish. Our intelligent ancestors might have discovered ways of streamlining that cumbersome process: e.g. instead of bringing each fish-eater to the river, ask each one to pick up a bowl and place it on the fisherman’s bowl. Then the bowls could be taken instead of the people, and the fisherman could give each bowl a fish, until there are no more empty bowls, then carry the laden bowls back.

What sort of brain mechanism would enable the first person who thought of doing that to realise, merely by *thinking* about it, that it *must* produce the same end result as taking all the people to the river? A non-mathematical (statistical reasoning) individual would need to be convinced by repetition that the likelihood of success is high. A mathematical mind would see the necessary truth. How?

Of course, we also find it obvious that there’s no need to take a collection of bowls or other physical objects to represent individual fish-eaters. We could have a number of blocks with marks on them, a block with one mark, a block with two marks, etc., and any one of a number of procedures for matching people to marks could be used to select a block with the right number of marks to be used for matching against fish.

Intelligent fishermen could understand that a collection of fish matching the marks would also match the people. How? Many people now find that obvious but realising that one-one correspondence is a transitive relation is a major intellectual achievement, crucial to abilities to use numbers. I doubt that any neuroscientist knows how brains support this realisation. We also know that it is not necessary to carry around a material numerosity indicator: we can memorise a sequence of names and use each name as a label for the numerosity of the sub-sequence up to that name, as demonstrated in (Sloman, 1978a, Chap8)¹⁷. A human-like intelligent machine would also have to be able to discover such strate-

¹⁷ <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap8>

gies, and understand why they work. This is totally different from achievements of systems that do pattern recognition. Or deep statistical learning. Perhaps studying intermediate competences in other animals will help us understand what evolution had to do to produce human mathematicians. Understanding mathematical induction over infinite sets is an even greater challenge. (This is all much deeper than learning to assign number names.)

Piaget's work showed that five- and six-year old children have trouble understanding consequences of transforming 1-1 correlations, e.g. by stretching one of two matched rows of objects (Piaget, 1952). When they do grasp the transitivity have they found a way to derive it from some set of logical axioms using explicit definitions? Or is there another way of grasping that if two collections A and B are in a 1-1 correspondence and B and C are, then A and C *must* also be, even if C is stretched out more in space?

I suspect that for most people this is more like an obvious topological theorem about patterns of connectivity in a graph rather than something proved by logic.

But why is it obvious to adults and not to 5 year olds? Anyone who thinks it is merely a probabilistic generalisation that has to be tested in a large number of cases has not understood the problem, or lacks the relevant mechanisms in normal human brains. Does any neuroscientist understand what brain mechanisms support discovery of such mathematical properties, or why they seem not to have developed before children are five or six years old (if Piaget used good tests). Much empirical research on number competences grossly over simplifies what needs to be explained, omitting the role of reasoning about 1-1 correspondences.

It would be possible to use logic to encode the transitivity theorem in a usable form in the mind of a robot, but it's not clear what would be required to mirror the developmental processes in a child, or our adult ancestors who first discovered these properties of 1-1 correspondences. They may have used a more general and powerful form of *relational* reasoning of which this theorem is a special case. The answer is not statistical (e.g. neural-net based) learning. Intelligent human-like machines would have to discover deep non-statistical structures of the sorts that Euclid and his precursors discovered.

The machines might not know what they are doing, like young children who make and use mathematical or grammatical discoveries. But they should have the meta-cognitive ability to become self-reflective and later make philosophical and mathematical discoveries. I suspect human mathematical understanding requires at least four layers of meta-cognition, each adding new capabilities, but will not defend that here. Perhaps robots with such abilities in a future century will discover how evolution produced brains with these capabilities (Sloman, 2013).

Close observation of human toddlers shows that before they can talk they are often able to reason about consequences of spatial processes, including a 17.5 month pre-verbal child apparently testing a sophisticated hypotheses about 3-D topology, namely: if a pencil can be pushed point-first through a hole in paper from one side of the sheet then there must be a continuous 3-D trajectory by which it can be made to go point first through the same hole from the other

side of the sheet:¹⁸. (I am not claiming that my words accurately describe her thoughts: but clearly her intention has that sort of complex structure even though she was incapable of saying any such thing in a spoken language. What sort of GL was she using? How could we implement that in a baby robot?)

Likewise, one does not need to be a professional mathematician to understand why when putting a sweater onto a child one should not start by inserting a hand into a sleeve, even if that is the right sleeve for that arm (Note 1). Records showing 100% failure in such attempts do not establish impossibility, since they provide no guarantee that the next experiment will also fail. Understanding impossibility requires non-statistical reasoning.

Generalising Gibson

James Gibson proposed that the main function of perception is not to provide information about what occupies various portions of 3-D space surrounding the perceiver, as most AI researchers and psychologists had previously assumed (e.g. (Clowes, 1971; Marr, 1982)) but rather to provide information about what the perceiver can and cannot do in the environment: i.e. information about positive and negative *affordances* - types of possibility.

Accordingly, many AI/Robotic researchers now design machines that learn to perform tasks, like lifting a cup or catching a ball by making many attempts and inferring probabilities of success of various actions in various circumstances.

But that kind of statistics-based knowledge cannot provide mathematical understanding of what is *impossible*, or what the *necessary* preconditions or consequences of certain spatial configurations and processes are. It cannot provide the kind of reasoning capabilities that led up to the great discoveries in geometry (and topology) (e.g. by Euclid and Archimedes) long before the development of modern logic and the axiomatic method. I suspect these mathematical abilities evolved out of abilities to perceive a variety of positive and negative affordances, abilities that are shared with other organisms (e.g. squirrels, crows, elephants, orangutans) which in humans are supplemented with several layers of metacognition (not all present at birth).

Spelling this out will require a theory of modal semantics that is appropriate to relatively simple concepts of possibility, impossibility and necessary connection, such as a child or intelligent animal may use (and thereby prevent time-wasting failed attempts).

What sort of modal semantics?

I don't think any of the forms of "possible world" semantics are appropriate to the reasoning of a child or animal that is in any case incapable of thinking about the *whole* of *this* world let alone sets of alternative possible worlds. Instead, the learner's modal semantics will have to be based on a grasp of ways in which properties and relationships in a *small portion* of the world can change and which combinations are possible or impossible. E.g. if two solid rings are linked it is impossible for them to become *unlinked* through any continuous form of motion

¹⁸ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#pencil>

or deformation – despite what seems to be happening on a clever magician’s stage. This form of modal semantics, concerned with possible rearrangements of a *portion* of the world rather than possible *whole* worlds was proposed in (Sloman, 1962). Barbara Vetter seems to share this viewpoint (Vetter, 2013). Another type of example is in the following figure.

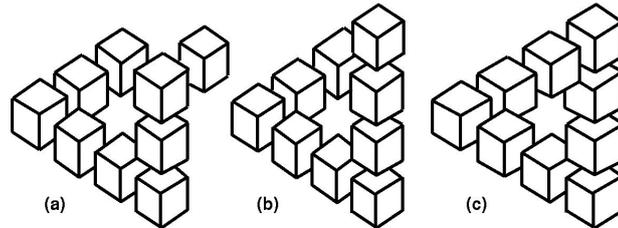


Fig. 1. Possible and impossible configurations of blocks.
(Swedish artist Oscar Reutersvard drew the impossible configuration in 1934)

What sort of visual mechanism is required to tell the difference between the possible and the impossible configurations. How did such mechanisms evolve? Which animals have them? How do they develop in humans? Can we easily give them to robots? How can a robot detect that what it sees depicted is impossible?¹⁹

A child could discover prime numbers by attempting to arrange groups of blocks into $N \times M$ regular arrays. It works for twelve blocks but adding or removing one makes the task impossible. Has any child ever discovered primeness in that way? Which robot will be the first to do that? (Pat Hayes once informed me that a frustrated conference receptionist trying to tidy uncollected name cards made that discovery without recognizing its significance. She thought her inability to make a rectangle in some cases was due to stupidity.)

The link to Turing

What might Alan Turing have worked on if he had not died two years after publishing his 1952 paper on the Chemical basis of morphogenesis? Perhaps the Meta-Morphogenesis (M-M) project: an attempt to identify significant transitions in types of information-processing capabilities produced by evolution, and products of evolution, between the earliest (proto-)life forms and current organisms, including changes that modify evolutionary mechanisms²⁰.

Conclusion

Natural selection is more a blind mathematician than a blind watchmaker: it

¹⁹ Richard Gregory demonstrated that a 3-D structure can be built that looks exactly like an impossible object, but only from a particular viewpoint, or line of sight.

²⁰ <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

discovers and uses “implicit theorems” about possible uses of physics, chemistry, topology, geometry, varieties of feedback control, symmetry, parametric polymorphism, and increasingly powerful cognitive and meta-cognitive mechanisms. Its proofs are implicit in evolutionary and developmental trajectories. So mathematics is not a human creation, as many believe.

The “blind mathematician” later produced at least one species with meta-cognitive mechanisms that allow individuals who have previously made “blind” mathematical discoveries (e.g. what I’ve called “toddler theorems”) to start noticing, discussing, disputing and building a theory unifying the discoveries.

Later still, meta-meta-(etc?)cognitive mechanisms allowed products of meta-cognition to be challenged, defended, organised, and communicated, eventually leading to collaborative advances, and documented discoveries and proofs, e.g. Euclid’s Elements (sadly no longer a standard part of the education of our brightest learners). Many forms of applied mathematics grew out of the results. Unfortunately, most of the pre-history is still unknown and may have to be based on intelligent guesswork and cross-species comparisons. Biologically inspired future AI research should provide clues as to currently unknown intermediate forms of biological intelligence. I don’t think current AI techniques are anywhere near this, especially statistics-based AI. Perhaps we’ll make more progress by trying to model less sophisticated reasoners, like Portia Spiders and intelligent toddlers (with their rich internal languages) than attempting to replicate Euclid directly.

Acknowledgements:

I am grateful for discussions with Jackie Chappell about animal intelligence, with Aviv Keren about mathematical cognition, and discussions about life, the universe, and everything with Birmingham colleagues and Alison Sloman.

References

- Chappell, J., & Sloman, A. (2007). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, 3(3), 211–239. Available from <http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717>
- Clowes, M. (1971). On seeing things. *Artificial Intelligence*, 2(1), 79–116.
- Furbach, U., & Schon, C. (Eds.). (2016). New York: To be published by Springer. (<http://ratiolog.uni-koblenz.de/proceedings2016.pdf>)
- Glasgow, J., Narayanan, H., & Chandrasekaran, B. (Eds.). (1995). *Diagrammatic reasoning: Computational and cognitive perspectives*. Cambridge, MA: MIT Press.
- Jablonka, E., & Lamb, M. J. (2005). *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge MA: MIT Press.
- Kant, I. (1781). *Critique of pure reason*. London: Macmillan. (Translated (1929) by Norman Kemp Smith)

- Lakatos, I. (1976). *Proofs and Refutations*. Cambridge, UK: Cambridge University Press.
- Marr, D. (1982). *Vision*. San Francisco: W.H. Freeman.
- Piaget, J. (1952). *The Child's Conception of Number*. London: Routledge & Kegan Paul.
- Rescorla, M. (2015). The computational theory of mind. In E. N. Zalta (Ed.), *The stanford encyclopedia of philosophy* (Winter 2015 ed.).
- Senghas, A. (2005). Language Emergence: Clues from a New Bedouin Sign Language. *Current Biology*, 15(12), R463–R465.
- Sloman, A. (1962). *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil Thesis)*.
(<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962>)
- Sloman, A. (1971). Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In *Proc 2nd ijcai* (pp. 209–226). London: William Kaufmann.
- Sloman, A. (1978a). *The computer revolution in philosophy*. Hassocks, Sussex: Harvester Press (and Humanities Press) Revised 2015.
(<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp>)
- Sloman, A. (1978b). What About Their Internal Languages? Commentary on three articles in *BBS Journal* 1978, 1 (4). *BBS*, 1(4), 515.
- Sloman, A. (1979). The primacy of non-communicative language. In M. MacCafferty & K. Gray (Eds.), *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979* (pp. 1–15). London: Aslib.
- Sloman, A. (2005). *Discussion note on the polyflap domain (to be explored by an 'altricial' robot)* (Research Note No. COSY-DP-0504). Computer Science, Univ. of Birmingham.
- Sloman, A. (2013). *Meta-Morphogenesis and Toddler Theorems: Case Studies*. School of Computer Science, The University of Birmingham. Available from <http://goo.gl/QgZU1g> (Online discussion note)
- Sloman, A. (2015). *What are the functions of vision? How did human language evolve?* (Online research presentation)
- Sloman, A., & Chappell, J. (2007). Computational Cognitive Epigenetics (Commentary on (Jablonka & Lamb, 2005)). *BBS*, 30(4), 375–6.
- Tarsitano, M. (2006, December). Route selection by a jumping spider (*Portia labiata*) during the locomotory phase of a detour. *Animal Behaviour*, 72, Issue 6, 1437–1442.
- Vetter, B. (2013, Aug). 'Can' without possible worlds: semantics for anti-Humeans. *Imprint Philosophers*, 13. (16)
- Weir, A. A. S., Chappell, J., & Kacelnik, A. (2002). Shaping of hooks in New Caledonian crows. *Science*, 297(9 August 2002), 981.
- Whitehead, A. N., & Russell, B. (1910–1913). *Principia Mathematica Vols I – III*. Cambridge: Cambridge University Press.