**DRAFT:Updated: Apr 25, 2016:DRAFT**
**Temporarily available at:**
**http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-ai-gaps.html**
**http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-ai-gaps.pdf**
**...**

# Natural Vision and Mathematics: Seeing Impossibilities

## (Draft workshop paper)

**Aaron Sloman** [1]
**School of Computer Science,**
**University of Birmingham, UK**
**http://www.cs.bham.ac.uk/~axs**

## Abstract

This paper is itself an extended abstract for one aspect of a large and complex project - the Turing-inspired investigation of evolution of forms of information processing: the Meta-Morphogenesis project. The full project investigates forms of information processing produced by evolution at various times since the beginning of life on earth, and the fundamental and evolved construction kits that helped to make them possible. This paper zooms in on a few features of human and animal information processing that seem to be relevant to the task of understanding the mechanisms that made possible the deep mathematical discoveries made by Euclid, Archimedes, and other ancient mathematicians, partly on the basis of conjectured mechanisms of spatial perception that were precursors of such mathematical abilities. These are mechanisms involved in perception of possibilities and constraints on possibilities, a type of affordance perception not explicitly discussed by Gibson, but suggested by extending his ideas. Current AI vision systems and reasoning systems do not seem to be capable of anything like the same achievements. A future product of this project might be a design for a "baby" robot that is able to "grow up" to become a mathematician 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 design of many kinds of intelligent machine able to perceive and interact with structures and processes in the environment. No early results are promised. One consequence is demonstration of the need to extend Dennett's taxonomy of types of mind to include Euclidean (or Archimedean) minds.

**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 the history of mathematics up to Euclid and Archimedes, for example, and also ignores kinds of mathematical competence that are a part of our everyday life, but mostly go unnoticed. When noticed, they are generally not classified as mathematical, because they don't use numbers or special symbols. Yet they are major challenges for AI, especially attempts to replicate human mathematical competences. This is not the same kind of AI project as designing useful mathematical theorem provers, proof checkers, etc.

I'll explore some aspects of the evolution and operation of human mathematical competences and begin, but not finish, discussion of requirements for replicating those competences in future robots.

There were great mathematicians long before the development and use of modern logic, formal languages, and formal systems. Such things are therefore not *necessary* for mathematics, though they are *part* of mathematics: a fairly recent part. What is necessary, I claim, for pre-verbal toddlers, intelligent non-human animals and future human-like robots, is a type (or types) of powerful *internal* language, produced by evolution, for thinking about shapes, structures, processes and valid and invalid inferences about them. Criteria for such languages are accumulating, but slowly.

I first tried and failed, nearly 40 years ago, starting with [Sloman 1978 2,Sloman 1979] (and most recently [Sloman 2015]) to persuade researchers that structured *internal* languages (for storing and using information) must have evolved *before* languages for communication. There would be nothing to communicate and no use for anything communicated, without pre-existing internal mechanisms for constructing, manipulating and using structured meanings. For the simplest organisms (viruses?) there may be no decisions and no uses of information, merely passive physical/chemical reactions. 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 sit studying 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-dependent 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 any of an enormous variety of possible scene structures and use them to create and then use an appropriate 3-D route plan.

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

Human and non-human abilities to perceive varied plants in natural environments, for example in woodlands or meadows, and either immediately or at a later time 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 such objects or object-parts change their relationships, or change their visual projections as the perceiver moves.

The ability to act on perceived structures, e.g. to bite and consume them, or carry 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:

Animal competences requiring construction and use of information structures of varying complexity are evolutionary precursors of human abilities to use grammars and semantic rules for a human language so that novel sentences can be understood in systematic ways (e.g. using a lexicon, syntactic structure, and compositional semantics) to express different more or less complex percepts or intentions or plans to solve practical problems. In particular, complex new information structures can be assembled and stored that later serve as information structures (e.g. plans, hypotheses) used in control of 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.

Many nest-building actions are neither random, nor fixed repetitive movements. They are guided 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 later. Nest building competences in corvids and weaver birds are examples. Human architects are more complicated.

One of the benefits of creating, perceiving, changing, manipulating, or using meaning structures (of varying complexity) is enabling a perceiver of a novel situation or object to take in its structure and reason hypothetically about some of its possible behaviours if acted on - 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.

The "polyflap" domain consisting of arbitrary 2D polygonal shapes each with a single (non-flat) fold forming a new 3D object, was proposed in [Sloman 2005] as an artificial, but challenging, example of an environment illustrating the cognitive requirements. It may still be too difficult for current robots. For example: any object resting on surfaces where it has a total of two contact points may rotate in either direction about the line joining the contact points. Noticing this should allow a polyflap user to work out the need for at least one more supporting surface on which a third part of the object can rest, in order to be stable. In the simple case all three points may be in the same horizontal plane: e.g. on a floor. But an intelligent agent that understands stability requirements (and anyone who reads and understands this sentence for the first time) can produce stability with three support points on different, non-co-planar surfaces, e.g. the tops of three pillars with different heights. In that case 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 goes through the interior of the triangle formed by the three support points then the structure will be stable. Can any current robot think like that? [2]

An animal, or machine, that understands such constraints will be able to reason in similar ways about *novel* configurations. These can be viewed as examples of types of perception of affordances consistent with the *spirit* of Gibson's theory even though Gibson himself did not (as far as I know) extend his theory to include use of geometrical or topological reasoning in deriving information about what would or would not be stable.

This contradicts a common interpretation of Gibson's affordances as features detected through statistical learning. (Of course *some* affordances can be statistics based.) 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 reasoning are very important, but still unexplained.

I conclude that for some of the capabilities of many intelligent non-human animals, as well as for humans, mechanisms evolved 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, nothing said here implies that the linguistic forms of internal languages have to be linear structures. In principle they could be trees, graphs, nets, map-like structures or types of structure we have not yet thought of.

The variety of types of animal that can perceive and act intelligently in relation to novel perceived environmental structures, natural or intentionally built, 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 languages with similar features used externally for communication [Sloman_Chappell 2007,Sloman 2015]. 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.

## Is there a circularity?

Prior to the development of computing systems with such abilities 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 demonstrate that there need not be an infinite regress: a very important discovery of the last seven or so decades. (I don’t have space for the (non-trivial) details here, but the workshop audience should not need them.) How brains achieve this is unknown, however. 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 and apes.

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. A footnote points to some examples[3].

The full repertoire of such 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, especially in humans, and to a lesser extent in other altricial species with "meta-configured" competences in the terminology of [Chappell_Sloman 2007].

Human sign languages are also richly structured but are not restricted to use of discrete temporal sequences of simple signs: usually movements of hands, head and parts of the face (e.g. eyes and mouth) 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. (J.L.Austin once wrote "Fact is richer than diction".)

## 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 have mechanisms, which expand in power over time as they are used, enabling them to *create* languages collaboratively. Normally they do this collaborative creation as a relatively powerless minority, so the creation produces results that look like imitative learning. The deaf children in Nicaragua showed that the process involves language creation rather than mere learning[4].

Although many details remain unspecified, I hope it’s clear that many familiar processes of perceiving, learning, intending, planning, plan execution, debugging faulty plans, etc. would be impossible if humans (and perhaps some other intelligent animals with related capabilities) did not have rich *internal* languages and language manipulation abilities. (GL competences.) There’s no other known way they could work! (Unless we are to believe in magic, or Wittgenstein’s sawdust in the skull.) For more on this see [Sloman 2015]. (There is a myth believed by some philosophers, cognitive scientists and others that structure-based "old fashioned" AI has failed. But the truth is that NO form of AI has "succeeded" as yet, except for powerful narrowly focused AI applications, and the newly fashionable versions are not necessarily closer to general success. I find them much shallower.[5])

There could not be any point developing mechanisms for communicating information, i.e. languages of the familiar type, if senders and recipients were not *already* information *users*, otherwise they would have nothing to communicate, and would have no way to change themselves when something has been understood. 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 many of the competences displayed by pre-verbal humans and other animals, and partly because they don't understand how the requirement does not lead to an infinite regress of internal languages.

Dennett (1995, and other publications) is an arch-opponent of this idea: his theory of consciousness argues, on the contrary, that consciousness followed evolution of mechanisms allowing languages previously used for external communication to be used internally for silent self-communication. That seems to imply that Portia spiders needed ancestors that discussed planned routes to capture prey before they evolved the ability to talk to themselves silently about the process in order to survey, plan, climb and feed unaided?

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, spider monkeys, apes, 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 meal.[6] Of course, 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 that there were explanatory gaps, 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*.[7] Thousands of people all around the world are still putting its discoveries to good use every day even if they have never read it.[8]

As a mathematics graduate student interacting with philosophers around 1958, my impression was that the philosopher whose claims about mathematics were closest to what I knew about the processes of *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 axiom.[9] 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.

First-hand experience of doing mathematics suggests 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 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 reasoning.[10]

Purely philosophical debates on these issues can be hard to resolve. So when Max Clowes[11] 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). 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 at IJCAI 1971 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 called metaphysically adequate if the world could have that form without contradicting the facts of the aspect of reality that interests us.*
> *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, for instance, 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 failing to realise 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, especially (among others) "analogical" representations that were not based on the kind of function/argument structure used by logical representations. Despite an explicit disclaimer in the paper it is often mis-reported as claiming that analogical representations are isomorphic with what they represent: which may be true in special cases, but is clearly false in general, 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 subtle challenge is provided by the discovery known to Archimedes that there is a simple and natural way of *extending* Euclidean geometry (the *neusis* construction) which makes it easy to trisect an arbitrary angle, as demonstrated here:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html[12]

I don’t think much is known about that sort of discovery process and 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*. It is also not a case of derivation 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 axiom:

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

I don’t know of anyone in AI who has tried to implement abilities to discover Euclidean geometry, including topological reasoning, or its various 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.)

It’s not only competences of adult human mathematicians that have not yet been replicated. Many intelligent animals, such as squirrels, nest building birds, elephants and even octopuses have abilities to perform spatial manipulation 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.

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 geometricians 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. I suspect that has only limited uses and a deeper advance can come from extending techniques for reasoning about possibilities, impossibilities and changing topological relationships and the use of partial orderings (of distance, size, orientation, curvature, slope, containment, etc.) as suggested in
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/changing-affordances.html`
I'll return to this topic below.

## 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 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 count 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, by *thinking* about it, that it *must* produce the same end result as taking all the people to the river? A non-mathematical 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. 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 1978 1,Chap8]. A human-like intelligent machine would also have to be able to discover such strategies, and understand why they work. This is totally different from achievements of systems that do pattern recognition. Perhaps studying intermediate competences in other animals will help us understand what evolution had to do to produce human mathematicians. (This is 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 (unless Piaget asked his subjects the wrong questions).[13]

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 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:
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html#pencil`.
(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. 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* consequences of certain spatial configurations and processes are. It cannot provide understanding of 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 I think the kind of 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 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 figure: 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?[14]

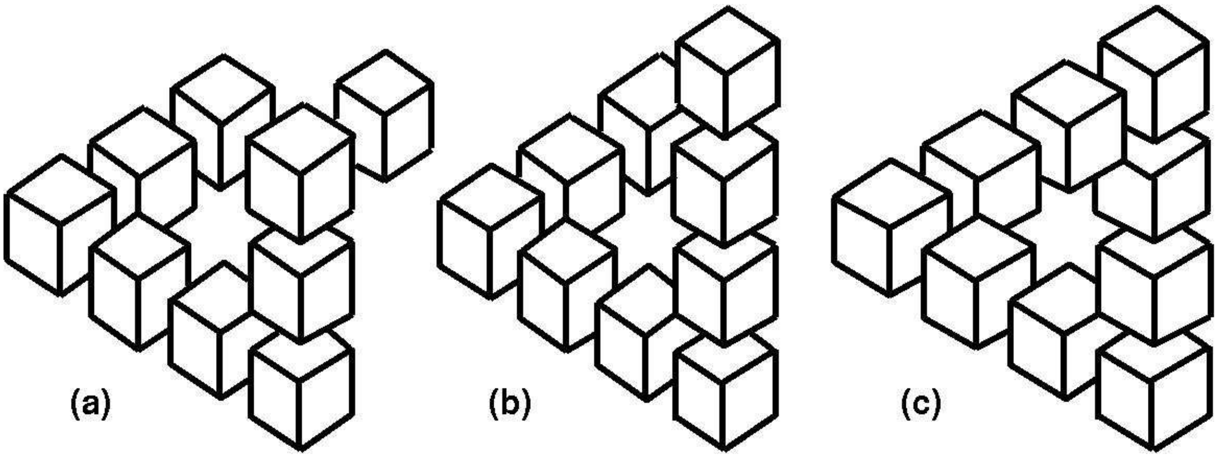


Figure 1: Possible and impossible configurations of blocks.
*(Swedish artist Oscar Reutersvard drew the impossible configuration in 1934)*

A child can in principle discover prime numbers by attempting to arrange different collections of blocks into NxM regular arrays. It works for twelve blocks but adding or removing one makes the task impossible. I don't know if any child ever has discovered primeness in that way, but it could happen. 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 failure on occasions to make a rectangle was due to her 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 mechanisms.
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html

## Conclusion

Natural selection is more a blind mathematician than a blind watchmaker: it 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, and the early forms of representation and reasoning are not necessarily similar to recently invented logical, algebraic, or probabilistic forms.
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 will provide clues

as to currently unknown intermediate forms of biological intelligence.

**Acknowledgements:**

This paper owes much to discussions with Jackie Chappell about animal intelligence, discussions with Aviv Keren about mathematical cognition, and discussions about life, the universe, and everything with Birmingham colleagues and Alison Sloman.

# References

[Chappell  Sloman 2007]
Chappell, J.  Sloman, A.  (2007). Natural and artificial meta-configured altricial information-processing systems. International Journal of Unconventional Computing, 3 (3), 211-239.

[Chomsky 1965]
Chomsky, N.  (1965). Aspects of the theory of syntax. Cambridge, MA: MIT Press.

[Clowes 1971]
Clowes, M.  (1971). On seeing things. Artificial Intelligence, 2 (1), 79-116.

[Dennett 1995]
Dennett, D.  (1995). Darwin's dangerous idea: Evolution and the meanings of life. London and New York: Penguin Press.

[Frege 1950]
Frege, G.  (1950). The Foundations of Arithmetic: a logico-mathematical enquiry into the concept of number. Oxford: B.H. Blackwell. ((Tr. J.L. Austin. Original 1884))

[Gibson 1979]
Gibson, J J.  (1979). The ecological approach to visual perception. Boston, MA: Houghton Mifflin.

[Glasgow, Narayanan,  Chandrasekaran; . 1995]
Glasgow, J., Narayanan, H.  Chandrasekaran, B. ().   (Eds.). (1995). Diagrammatic reasoning: Computational and cognitive perspectives. Cambridge, MA: MIT Press.

[ Jablonka  Lamb 2005]
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 1781]
Kant, I.  (1781). Critique of pure reason. London: Macmillan. (Translated (1929) by Norman Kemp Smith)

[Lakatos 1976]
Lakatos, I.  (1976). Proofs and Refutations. Cambridge, UK: Cambridge University Press.

[Marr 1982]
Marr, D. (1982). Vision. San Francisco: W.H.Freeman.

[ McCarthy  HayesMcCarthy  Hayes 1969]
McCarthy, J.  Hayes, P.  (1969). Some philosophical problems from the standpoint of AI. In B. Meltzer & D. Michie (Eds.), Machine Intelligence 4 (pp. 463-502). Edinburgh, Scotland: Edinburgh University Press.

[Piaget 1952]
Piaget, J.  (1952). The Child's Conception of Number. London: Routledge & Kegan Paul.

[Senghas 2005]
Senghas, A.  (2005). Language Emergence: Clues from a New Bedouin Sign Language. Current Biology, 15 (12), R463-R465.

[Sloman 1962]
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 1971]
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 1978 1]
Sloman, A.  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 1978 2]
Sloman, A.  (1978b). What About Their Internal Languages? Commentary on three articles in BBS Journal 1978, 1 (4). BBS , 1 (4), 515.

[Sloman 1979]
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 2005]
Sloman, A. Sloman, A. (2005, September). Discussion note on the polyflap domain (to be explored by an ‘altricial’ robot) (Research Note No. COSY-DP-0504). Birmingham, UK: School of Computer Science, University of Birmingham. Available from `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps`
[Sloman 2013]
Sloman, A. (2013). Meta-Morphogenesis and Toddler Theorems: Case Studies. School of Computer Science, The University of Birmingham. (Online discussion note) Available from `http://goo.gl/QgZU1g`
[Sloman 2015]
Sloman, A. (2015). What are the functions of vision? How did human language evolve? (Online research presentation) http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111
[Sloman Chappell 2007]
Sloman, A. Chappell, J. (2007). Computational Cognitive Epigenetics (Commentary on (Jablonka & Lamb, 2005)). BBS , 30 (4), 375-6.
[Tarsitano 2006]
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 2013]
Vetter, B. Vetter, B. (2013, Aug). ‘Can’ without possible worlds: semantics for anti-Humeans. Imprint Philosophers, 13 . (16)
[Weir, Chappell, Kacelnik . 2002]
Weir, A A S., Chappell, J. Kacelnik, A. (2002). Shaping of hooks in New Caledonian crows. Science, 297 (9 August 2002), 981.
[Whitehead Russell 1910-1913]
Whitehead, A N. Russell, B. (1910-1913). Principia Mathematica Vols I - III. Cambridge: Cambridge University Press.

---

## Footnotes:

[1] This is a snapshot of part of the Turing-inspired Meta-Morphogenesis project.

[2] I did not notice this "Polyflap stability theorem" until I tried to think of an example. I did not need to do any experiments and collect statistics to recognize its truth (given familiar facts about gravity). Do you?

[3] `www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html`

[4] This video is an eye-opener: `https://www.youtube.com/watch?v=pjtioIFuNf8`

[5] `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/chewing-test.html`

[6] `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision/plants` presents a botanical challenge for vision researchers.

[7] 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.

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

[9] `http://web.mnstate.edu/peil/geometry/C2EuclidNonEuclid/8euclidnoneuclid.htm`

[10] My 1962 DPhil thesis [Sloman 1962] presented Kant's ideas, before I had heard about AI. http://www.cs.bham.ac.uk/research/projects/cogaff/thesis/new

[11] http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-clowestribute.html
[12] I was unaware of this until I found the Wikipedia article in 2015: https://en.wikipedia.org/wiki/Angle_trisection#With_a_marked_ruler
[13] Much empirical research on number competences grossly over simplifies what needs to be explained, omitting the role of reasoning about 1-1 correspondences.
[14] 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.

---

File translated from TEX by TTH, version 3.85.
On 25 Apr 2016, 00:15.

---

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