(NB: Work in progress. Please save a link not a copy. Updated December 4, 2012)

Biological, computational and robotic connections

with Kant's theory of mathematical knowledge

Invited talk for ECAI 2012 Special "anniversary" session
Wednesday 29th August 2012
`http://www2.lirmm.fr/ecai2012/`

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

This will be added to my talks directory (PDF):
`http://tinyurl.com/BhamCog/talks/#anniv`

Draft paper available

There's a draft preprint paper on this topic available on the CogAff web site.

`http://www.cs.bham.ac.uk/research/projects/cogaff/12.html#1205`

The style is idiosyncratic and personal – because of the subject matter.

A revised version is due in November.

None of my papers is ever finished. I go on revising them after publication.

Anything of mine in print is out of date

I expect that will become the norm
– as with software, car designs, musical compositions, plays, proofs (?), etc.

Abstract / Apology

In my research I meander through various disciplines, using fragments of AI that I regard as relevant to understanding natural and artificial intelligence, willing to learn from anyone.

As a result, all my knowledge of work in particular sub-fields of AI is very patchy, and rarely up to date. This makes me unfit to write the history of European collaboration on some area of AI research as originally intended for this panel session.

However, by interpreting the topic rather loosely, I can (with permission from the event organisers) regard some European philosophers who were interested in Philosophy of mathematics as early AI researchers from whom I learnt much, such as Kant and Frege. Hume's work is also relevant.

Moreover, more recent work by neuro-developmental psychologist Annette Karmiloff-Smith, begun in Geneva with Piaget then developed independently, helps to identify important challenges for AI (and theoretical neuroscience), that also connect with philosophy of mathematics and the **future** of AI and robotics, rather than the history.

I'll present an idiosyncratic, personal, survey of a subset of AI stretching back in time, and deep into other disciplines, including philosophy, psychology and biology, and possibly also deep into the future, linked by problems of explaining human mathematical competences.

The unavoidable risk is that someone in AI has done very relevant work on mathematical discovery and reasoning, of which I am unaware.

I'll be happy to be informed, and will extend these slides if appropriate.

Meta-morphogenesis

This is part of the meta-morphogenesis project, as explained here:
http://tinyurl.com/BhamCog/misc/meta-morphogenesis.html

The name was suggested by Turing's 1952 paper on morphogenesis: He described processes that could produce large scale physical structures and patterns from small scale molecular interactions (e.g. diffusion and reaction of two chemicals, during development).

By analogy, I sketch an attempt to understand how biological evolution, and its products, including learning, development, and formation of social groups or ecosystems, can start with relatively simple forms of information processing, or even with just a supply of matter and energy, and, over time, produce extremely complex, powerful and varied forms of information processing.

Conjecture: Some products of mechanisms of morphogenesis of information processing turned out able to modify older mechanisms or introduce new mechanisms, so that the space of accessible products (some previously only theoretically possible) is constantly being expanded, and the processes of change and development are accelerated.

The project is not new – I've been working on it since I first learnt about AI in 1969, and probably many others have also, but without using the label "meta-morphogenesis".

I have a hunch that if Turing had lived longer, he might have developed his work in directions sketched below. Likewise Kenneth Craik (1943).

But those two would have made a lot more progress by now than I have!

Mathematical discovery as a biological phenomenon

It is clear that human mathematical capabilities depend on products of biological evolution.
What mechanisms were used by early mathematicians who discovered and proved mathematical results leading to Euclidean geometry?

Examples (expanded later):

- Moving a vertex of a triangle along a line parallel to the opposite side cannot change the area of the triangle;
- a three-sided closed polygon must have three angles, and the angles must add up to a straight line (180 degrees – see Mary Pardoe's proof, below);
- the result of counting a set of enduring, distinct objects does not depend on the order of counting.
- seven identical wooden cubes cannot be arranged to form a NxM rectangle with $N > 1$ and $M > 1$.

There is something about the information-processing of humans, and possibly some other species, that enables them to discover, and use, those and similar facts, which are immune from refutation by future experiments or observations. (Unlike statistical discoveries.)

Kant labelled them "necessary", "apriori" and "synthetic" – incapable of having exceptions, discoverable non-empirically, and able to add to our knowledge. (Sloman, 1965)

How can animals or machines make discoveries with such features?

And how can such discoveries have the characteristic of being "necessarily true", i.e. nothing can turn up in the world that contradicts them?

The philosophers I met around 1960 in Oxford all seemed to think that Hume's view of knowledge was correct. Experience of doing mathematics proved them wrong. How?

David Hume

David Hume apparently believed there were only two kinds of knowledge

– Empirical knowledge

Checking such knowledge requires experiment with and observation of the actual contents of the world, and empirical generalisations can never be conclusively proved, though they can be refuted by finding a counter-example.

They cannot be known to be true without exception: e.g. we cannot prove there are no circumstances in which an unsupported apple moves upwards not downwards.

Even alleged particular facts may be based on flawed observation or measurement.

– "Analytic" knowledge

Such knowledge concerns defining relations between our concepts, and their logical consequences.

Analytic propositions are true by definition of the terms involved: they are (allegedly) trivial kinds of knowledge that provide no real content about anything.

Examples:

"All bachelors are unmarried"
"All triangles have three angles"
"Every bachelor uncle has a brother or sister"

According to Hume, all writings that are not about these two kinds of knowledge are nothing but "sophistry and illusion" and should be consigned to the flames.

He was mainly thinking of theology and deep sounding but empty metaphysical philosophy.

When Immanuel Kant read Hume, he disagreed, though he credited Hume with waking him from his "dogmatic slumbers". (Kant, 1781)

For a more detailed account of these distinctions see (Sloman, 1965)

Kant on Mathematics

When I started reading Kant's Critique of Pure Reason (1781), I felt that he understood the nature of mathematical knowledge better than the other philosophers I had encountered (alive or dead) – insofar as his ideas matched my own experience of actually doing mathematics.

He claimed that mathematical knowledge was both synthetic (non-analytic) and a priori (non-empirical), and expressed necessary (non-contingent) truths, though such knowledge is not innate, since newborn babies lack the concepts required, as well as the ability to make mathematical discoveries).

Mathematicians and the methods they use are not infallible, since they can make mistakes, as Imre Lakatos showed in detail in (Lakatos, 1976)).

Unsurprisingly, Kant, writing around 1781, had no precise ideas about the transitions required for a neonate to develop mathematical competences, though it could not be merely a matter of discovering reliable correlations: nor the type of statistical learning that now, unfortunately, dominates much of AI.

There are also deep differences between:

– what human mathematicians have been doing for many centuries, and
– machines using powerful axioms and rules of inference, to generate and check proofs.

It became clear to me around 1970 that philosophy of mathematics needs AI to provide deep theories, and AI needs philosophy of mathematics to help us specify requirements for machines to be describable as doing mathematics, or making, understanding and using mathematical discoveries – like the prehistoric developers of Euclidean geometry.

Doing mathematics and understanding affordances

At IJCAI 1971 (Sloman, 1971) I criticised the logicist approach to AI presented and defended by McCarthy and Hayes (McCarthy & Hayes, 1969) in which they attempted to show that predicate logic (with fluents) was an adequate form of representation for intelligent agents.

Echoing Chomsky's notion of "adequacy" in (Chomsky, 1965), they claimed that logic had

- metaphysical adequacy (it could express anything that is true)
- epistemic adequacy (it can express anything an intelligent machine needs to know or represent)
- heuristic adequacy (it meets all requirements for tractable means of reasoning).

(This is a compressed summary: read the original 1969 paper.)

My claim was that the ability shared by humans (and in limited ways also some other animals, and very young children) to perceive and reason about possibilities for change in spatial configurations could be the basis of forms of reasoning that are sometimes more effective than logic-based forms of reasoning.

Later I learnt about James Gibson's ideas about perception in animals being concerned not merely with acquiring information about what actually exists in the environment, but also information affordances, e.g. information about which actions are possible, which impossible, and what the consequences of actions would be.

That led to the conjecture that the evolved ability to perceive and make use of affordances in the environment was an essential precursor to the ability to discover the concepts and theorems that were later systematically organised in Euclidean geometry.

This required extending Gibson's ideas about affordances in various ways:
`http://tinyurl.com/CogTalks/#gibson`
What's vision for, and how does it work? From Marr (and earlier) to Gibson and Beyond

Has AI made progress?
Yes lots – mostly engineering
Not so much in science and philosophy:
Deep aspects of natural intelligence remain unexplained

Has research in AI contributed to our understanding of mathematical learning, discovery and reasoning?

I shall try to show that this question has different interpretations, leading to different answers.

At first sight the question may seem pointless, given all the progress in automated theorem proving and the variety of computer-based mathematical tools used all round the world. E.g. see

http://en.wikipedia.org/wiki/Automatedtheoremproving
http://en.wikipedia.org/wiki/Automatedreasoning
http://en.wikipedia.org/wiki/Computationalgeometry
http://en.wikipedia.org/wiki/Scientificcomputing

See also Alan Bundy's talk in this session. [REF]

AI as science has deep connections with mathematics

Turing's original motivation: to design a machine with the minimal capabilities required for mathematical competences.

Much AI has used mathematics.

Some AI has extended mathematics.

The challenge of replicating human mathematical competences is deep and difficult – except in limited special cases.

(E.g. numerical calculations)

But progress in replicating human reasoning has been very slow, as I'll try to explain below.

The difference between discovering and using empirical correlations in controlling actions, and the discovery and use of mathematical truths (e.g. truths of Euclidean geometry) is closely related to the different requirements for online intelligence and offline intelligence.

The former involves use of feedback and continually changing control decisions to produce some action.

The latter involves understanding possibilities for and constraints on actions, without necessarily doing anything.

The current enthusiasm in some circles for "embodied cognition" or "enactivism" is largely based on a failure to understand that distinction, and its biological importance.

Gaps in AI

There has been much progress in automated theorem proving and many computer-based mathematical tools are used all round the world every day
– based on AI research, whether labelled as such or not.

Yet no existing AI system comes close to replicating the mechanisms involved in kinds of learning about the environment that develop into mathematical learning in young children, and must have contributed to the beginnings of human mathematical awareness.

Characterising what's missing is a subtle and difficult task, discussed below.

In the distant past, such biological capabilities must have led to the development of Euclidean geometry, and also use of numbers and discovery of properties of numbers – long before there were mathematics teachers. (Sloman, 2010)

Conjecture:
Prerequisites for making "proto-mathematical" discoveries include use of information structures with structural variability, generative mechanisms, and compositional semantics.

As argued in `http://tinyurl.com/BhamCog/talks/#glang` and (Sloman, 1979).

These are also features of forms of representation and mechanisms needed for:
– perception of complex structures, processes, and affordances,
– motive formation,
– forming and executing plans
– encoding various results of learning

underlying many capabilities developed by pre-verbal children and other animals.

What are the forms of representation and mechanisms used?
They cannot be directly inspected, but we can collect requirements.

Toddler Theorems and Proto-Mathematics

There are many types of "Toddler theorem" that young children can discover, through play and experiment with objects in the environment – though not all individuals can discover the same ones.

A growing collection of examples can be found here:

`http://tinyurl.com/BhamCog/misc/toddler-theorems.html`

`http://tinyurl.com/BhamCog/talks/#toddler`

Examples welcome: developmental psychologists mostly do not investigate such matters, partly because processes of discovery and learning are highly individual and trajectories very variable.

As a result, the topic requires something different from performing experiments providing data to be fed into statistical packages – often wrongly assumed to be a requirement for doing science.

We need to extend Gibson's theory of affordances

See: (Gibson, 1979), (Sloman, 2009), and `http://tinyurl.com/BhamCog/talks/#gibson`

There is reason to believe that various subsets of human "proto-mathematical" competences can be found in many other species, including squirrels, primates, nest-building birds, elephants, cetaceans (whales, dolphins, etc.)

They all seem to be able to make use of abstract features of spatio-temporal structures in order to select actions that will achieve desired results, for instance adding a twig to a partially built nest, or pushing down a tree to get at food that's out of reach near its top.

Squirrels are notoriously clever at defeating "squirrel-proof" bird-feeders.

Search for videos on Youtube, etc.

Informal human mathematics

Animal competences involving perception and use of affordances seem to be precursors of competences required for making discoveries about geometry.

I'll present some relatively simple examples of problems in Euclidean geometry and describe natural human ways of solving them that do not seem to map onto any known mechanised methods of reasoning, though there are different methods that give a superficial appearance of success.

In particular, the impressive work on mechanised theorem-proving over several decades produces machines that exceed the competences of most humans – but what they do is different from the geometric reasoning described below.

There are physics simulators and game engines that are able to predict in great detail consequences of events in a configuration of physical objects – but that is also different from mathematical reasoning in ways that I'll try to explain.

There are statistical learning systems that are able to infer probabilities of consequences of events, and can use those in controlling machinery – but the discoveries that led to the development of Euclidean mathematics were concerned with possibilities and impossibilities, not probabilities, as explained below.

There is a kind of non-probabilistic necessity in results of mathematical discovery, but that does not imply that mathematicians are infallible: mistakes can be made, and are made, then later discovered and usually remedied. (Lakatos, 1976)

Theorems about stretching a triangle

In the next few slides we'll discuss geometrical proofs concerning triangles.

- The median stretch theorem (MST):

  If a triangle is stretched by moving a vertex outward along a median the area must increase.

- The perpendicular stretch theorem (PST):

  If a triangle is stretched by moving a vertex in a perpendicular direction away from the opposite side the area must increase.

- The triangle area theorem (TAT):

  The area of a triangle is equal to half the area of a rectangle with the same base and height.

  From this it follows that moving a vertex parallel to the opposite side cannot increase or decrease the area.

MST is very easy to prove using diagrams.

PST has a sub-case that is easy to prove in the same way, and a sub-case that is more difficult to prove.

The difficult sub-case of PST can be proved by proving TAT.

I'll now use these theorems (and others) to illustrate biological forms of reasoning.

Reasoning about areas and containment

What happens to the area of a triangle if you slide a vertex of the triangle along a median through it?

A median is a straight line through the midpoint of one of the sides and the opposite vertex, as shown in the two diagrams on the right.

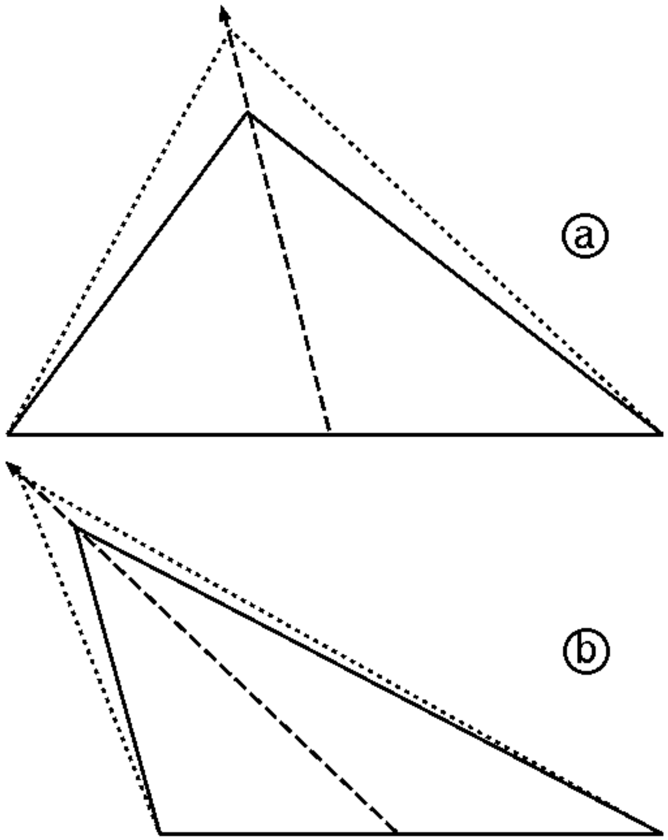

Question:

What happens to the area if a vertex moves along the median in the direction away from the opposite side?

Two cases are shown in figures (a) and (b) opposite.
In both cases the new triangle formed using the two new dotted sides, with the old sides unchanged, must have a larger area than the original triangle because the original (with solid lines) is entirely enclosed by the new one.

What will happen if the vertex moves along the median in the reverse direction: towards the opposite side?

How can you be sure the answers will remain
the same for **any** triangle with one vertex moved along the median through it?

Humans can look at a configuration and detect a constraint on possible changes in the configuration that ensures that no change of size, angles, orientation, location, colour, temperature, etc. can affect the constraint.

As far as I know no current AI program can do this with geometric structures.

More on reasoning about areas

Suppose we change the previous example and move the vertex on a line **perpendicular** to the opposite side, rather than along **the median**.

From visual inspection of the top figure on the right, it may seem that motion of the vertex along the perpendicular away from the opposite side must always increase the area of the triangle.

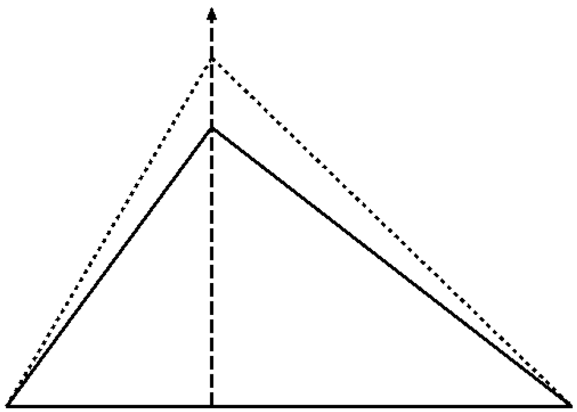

But what happens if, as in (b) on the previous slide, the vertex is no longer directly above the opposite side?

An example is shown below on the right: where motion upward on the perpendicular will not produce a new triangle that encloses the old triangle: because raising the vertex causes the lower (dotted) side of the new triangle to go **through** the old triangle instead of remaining **outside** it, as happens with motion on the median.

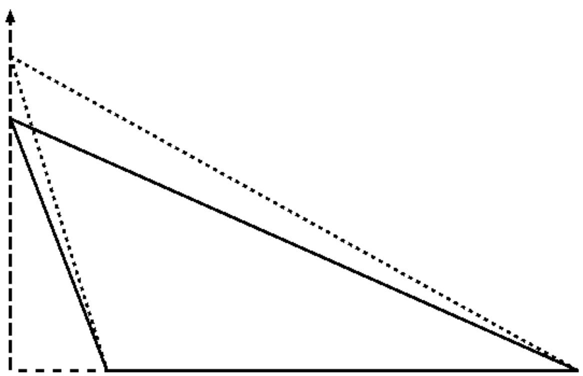

NB: we cannot argue that the area must increase because area of a triangle is $\frac{1}{2}$(base$\times$height), without first proving that formula.

Some mathematicians will be tempted to produce a proof by transforming the problem into a numerical/algebraic problem using the correspondence between geometry and arithmetic discovered by Descartes.

But that shifts the problem to: how can we give machines the ability to understand the Cartesian correspondence?

Triangles and rectangles

At some point an ancient mathematician may have discovered that any triangle can be enclosed in a rectangle, with all vertices of the triangle on sides of the rectangle, and at least two vertices of the triangle coinciding with corners of the rectangle.

It follows that at least one of the sides of the triangle lies on a side of the rectangle.

As the figure shows, this can be achieved in various different ways, depending on the shape of the triangle.

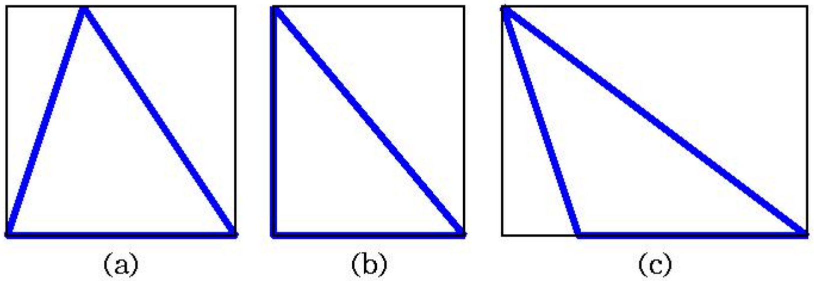

By "hallucinating" additional lines onto such figures it is possible to derive a fixed relationship between the area of the triangle and the area of the rectangle, or possibly another, smaller rectangle.

(Which case requires a new smaller rectangle to be constructed?)

That relationship proves the "conjecture" on the previous slide that motion on the perpendicular away from the opposite side always increases the area of the triangle.

Details are left as an exercises for the reader.

If you investigate this try to identify the possibilities that you detect, how you detect them, and the invariants you observe for different possibilities, and how you can be sure of them.

A hint: Handling case (c) requires subtraction of areas, unlike cases (a) and (b).

The need for case analysis

The difference between the example of motion of a vertex of a triangle along a median and motion on a perpendicular to the opposite side indicates a reason why mathematical reasoning, though non-empirical is not infallible.

Some problems can be solved by taking a special case and finding an invariant that can be seen to apply to all possible cases.

However, other problems may require subdivision into cases that have to be dealt with by slightly different reasoning: and sometimes human mathematicians fail to notice that there are additional cases, and wrongly conclude that a conclusive proof has been found.

Imre Lakatos showed in (Lakatos, 1976) that there are examples in the history of mathematics (specifically Euler's formula for polyhedral solids: $V - E + F = 2$) where not all cases were discovered when a proof was originally presented.

New cases were discovered on several occasions requiring previous proofs or formulations of the theorem to be modified.

Lakatos concluded that mathematics is *quasi-empirical* insofar as there are partial analogies between searching for empirical evidence against a scientific theory and searching for special cases that refute a mathematical generalisation.

Despite all this there are cases where a fairly simple form of inspection of examples reveals an invariant that covers all cases, so that a finite observation proves something about infinitely many different cases, e.g. MST: the median stretch theorem.

Another example is an unusual proof of the triangle sum theorem, on the next slide.

Taming infinities (continua)

A crucial feature of many of the discoveries in geometry is that they are applicable to an infinite range of cases, even if the discovery was initially made by thinking about one special case: what is discovered is an invariant common to a range of possibilities.

Example:

The triangle sum theorem

Why must interior angles of a planar triangle always add up to a straight line?

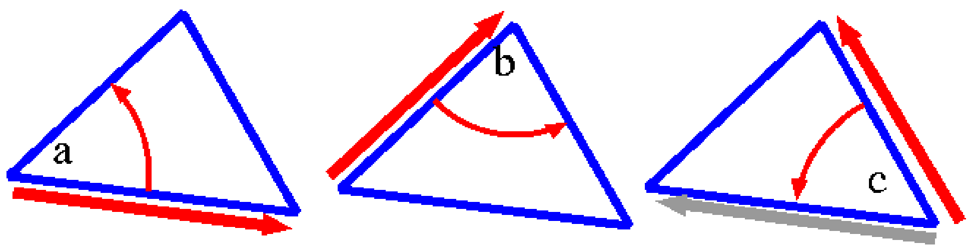

We can answer that by considering structural relationships within one triangle.

Mary Pardoe's proof is above:
The arrow, initially on the base, can be rotated through angles a, b and c.

It ends up on the starting line, but pointing in the opposite direction, having done half a revolution, so the internal angles must sum to 180^o

How do you know that the infinitely many possible changes in size shape, orientation or location of the triangle cannot affect the result?

Note: this proof breaks down on a sphere, e.g. if the triangle has its base on the equator of the earth and top vertex on the north pole, the three angles could each be 90^o
But the standard proof using Euclid's parallel axiom also breaks down in that case.

See also (Asperti, 2012) and

`http://tinyurl.com/BhamCog/misc/p-geometry.html`

`http://tinyurl.com/BhamCog/misc/triangle-theorem.html`

Using hallucinated structure in a proof

In the previous cases a theorem about a structure or process was proved by drawing additional lines.

In fact the lines do not need to be drawn: they can be imagined, or hallucinated.

It is often the case that a mathematical proof about some structure depends on the ability to perceive the possibility of an extension to that structure, i.e. the ability to hallucinate new structure.

An interesting example is discovering the concept of a prime number by considering changeable geometric configurations, presented in the next slide.

(Post-??)Toddler theorems about prime numbers

Playing with a collection of cubes of the same size you can try forming patterns with them.

Sometimes you'll fail: Why?

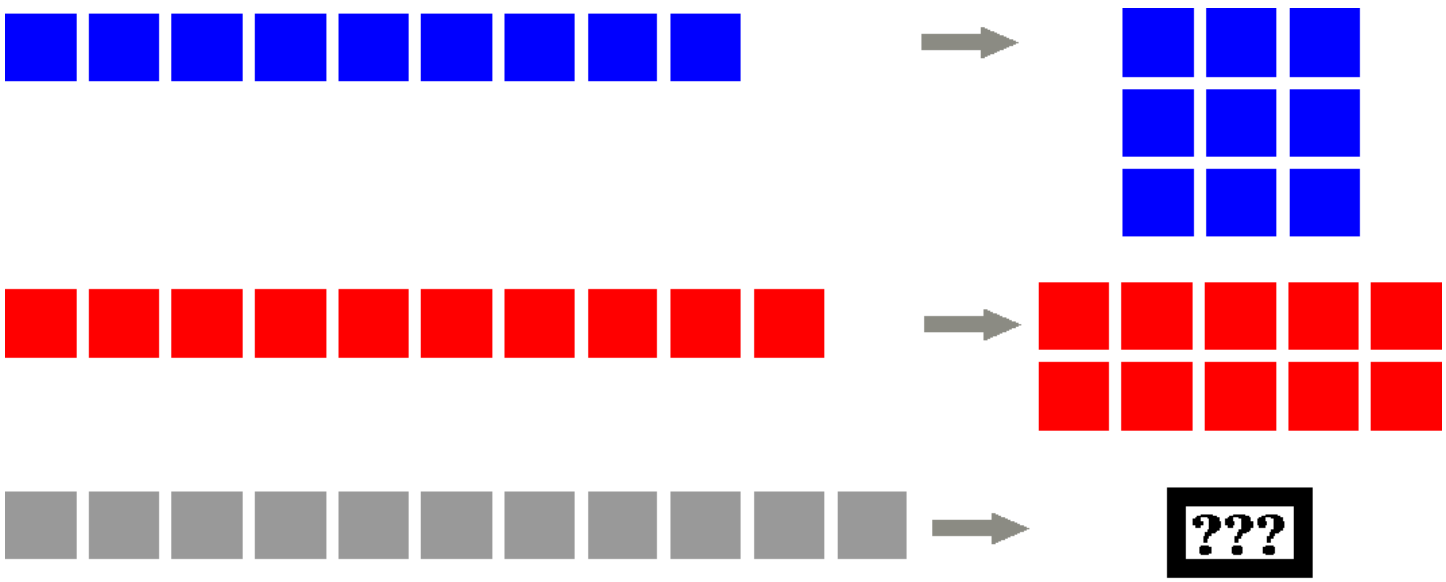

How can a child become convinced that something is impossible?

There are very many more examples.

NB: don't confuse learning about the 'natural numbers' with learning about numerosity of perceived collections – a far shallower competence.

Gabriella Daroćzy has started investigating what young children and also adults do when presented with these challenges.

Proving that something is impossible

Proving that no shorter configuration of the thirteen blocks can form a rectangle is difficult because it involves consideration of an infinite variety of possible configurations.

One way of doing this is to hallucinate a rectangular grid and exhaustively try different assignments of blocks to grid locations, to see if one produces a rectangle.

If the grid size is limited (e.g. by making it a bounded square grid) then the impossibility can be established by exhaustive systematic trial.

That still leaves the need to understand that that grid is representative of all possible grids containing the given number of blocks: another kind of abstraction covering infinitely many cases.

As far as I know, no current AI system can do that sort of reasoning.

For some steps in this direction see (Jamnik, Bundy, & Green, 1999)

Some possible payoffs

If we can find out how to model or replicate these biological forms of reasoning the result may be:

- new insights into the variety of possible forms of mathematical representation and reasoning;

  Especially cases involving structures and processes with a mixture of continuous and discrete change, where some of the invariants and implications can be "read off" special cases in ways that immediately generalise to a continuum of sub-cases. (For some comments see (Asperti, 2012))

- new ideas for other types of reasoning, e.g. action planning;
- a defence of a modern variant of Kant's philosophy of mathematics;
- far more versatile and creative machines;
- new insights into human reasoning capabilities, and how they develop

  (E.g. important alternatives to reasoning with probabilities, as discussed in (Sloman, 2007));

- new ways of thinking about the question whether all forms of reason can be replicated in a Turing machine, or whether there are additional powerful forms of discovery and reasoning, used by human mathematicians and others.

However, I suspect that achieving this understanding will require us to investigate far more transitions in biological information programming between the earliest, simplest organisms and our more recent ancestors, to help us understand the many layers of processing involved even in apparently simple thinking.

A growing list of examples of transitions in information processing is here (contributions welcome):

`http://tinyurl.com/CogMisc/evolution-info-transitions.html`

AI as science has deep implications for Biology

Researchers have found overlaps between AI and work in developmental psychology and animal cognition.

But there are two directions of influence:

Biologically-Inspired AI (BI-AI)
vs
AI-Inspired Biology (AI-IB)

– the former is generally shallow)

There have been several workshops and conferences on AIIB, not all using the label, in the last few years.
I suspect AIIB work will grow.

E.g. explaining squirrel intelligence is a challenge.
The picture shows a squirrel raiding a “squirrel- proof” bird feeder held by suckers high on a patio door.
The squirrel managed to climb up the the plastic-covered door frame just visible on left, then launch itself sideways across the glass, landing on the tray with nuts – a remarkable piece of creative intelligence, and ballistic control.

Grey squirrels defeat many bird-feeders sold as “squirrel-proof”.

They seem to be able to reason about what to do in advance of doing it, even in novel situations, requiring a primitive form of “theorem proving” about what can work? Robotic understanding of affordances is currently far inferior to animal understanding.
See (Gibson, 1979) and `http://tinyurl.com/BhamCog/talks/#gibson`

Craik on Animal Reasoning

Kenneth Craik (1943) suggested that in order not to have to discover empirically that certain actions would be harmful, or fatal, some animal species had evolved the ability to build models of novel situations and "run" them to work out consequences of possible actions.

I don't know whether he understood the importance of perception of and reasoning about affordances, later described by Gibson (Gibson, 1979) and generalised in `http://tinyurl.com/BhamCog/talks/#gibson`.

Or whether he understood the requirement to reason not only about very specific possibilities, like simulations running forward in a game engine, but also to use abstractions that covered classes of cases.

E.g. the squirrel in the previous photograph should not have to simulate in advance all the sensory and motor signals and their effects in order to select the climbing strategy: it would have to cope with an explosive continuum (or collections of continua) of possibilities.

The need for abstraction and chunking (discretisation) of affordances and possible consequences of actions, allowing finite case analysis to cover infinite collections of possibilities (e.g. continuously variable movements, or a continuum of possible target locations for rock in a cave), may have been major precursors of mathematical reasoning,

Craik worked before video cameras with frame-grabbers were common-place, so he did not fall into the trap of assuming vision starts from a rectangular grid of measurements with a well defined metric for grid coordinates – an assumption that blinkers many vision researchers.

Problems with AI/Cognitive science/Robotics "fads"

The need to be able to perceive, learn about, think about and make use of abstract possibilities, and branching sets of possibilities in situations requiring planning, explaining and problem solving, is often not noticed, or even ignored.

For example those human/animal competences are often ignored in research on "enactivist" or "embodied" cognition, which focuses only on online intelligence, the ability to produce and control detailed actions tailored to details of the environment.

This mostly ignores "offline" creative intelligence, involving thinking not just about what exists or is happening, but about what's possible, or could not happen: (Sloman, 1996, 2006, 2009)

A more subtle requirement is the ability to see commonalities or constraints in a space of possibilities, e.g. as discussed in:

`http://tinyurl.com/CogMisc/triangle-theorem.html`

`http://tinyurl.com/CogMisc/toddler-theorems.html`

Work on statistical/probabilistic AI tends to ignore the need to explain how mathematical understanding can arise, or how non-correlational causal understanding works, e.g. why do two meshed gear wheels with fixed axles have to rotate in opposite directions.

Learning how to prove things in school maths

Learning Euclidean geometry and developing skills at proving theorems and solving problems used to be a core part of the education of bright children.

Unfortunately educational fashions change and many children learn little or no mathematics, and even if they learn about theorems may merely learn to memorise them, and perhaps to use them.

How many children now encounter, let alone discover proofs for:

- The triangle sum theorem
- The formula for the area of a triangle
- Pythagoras theorem
- Rolle's theorem

A good mathematical education would include both:
– continuous mathematics
– discrete mathematics

And kinds of mathematics involving both continuity and discreteness, e.g. analysis/calculus

Euclidean geometry imposed rigidity and infinite precision on a continuous mush – HOW??

19th Century mathematics

In the 19th century there was a huge amount of progress in understanding notions like continuity and how to think about the continuum.

Eventually it became clear that if you start with just the natural numbers
0, 1, 2, 3, ...
you can define first the integers as sets of pairs of natural numbers
...-3, -2, -1, 0, +1, +2, +3 ...

Then sets of pairs of integers can be used to define rationals
not easy to represent as above because they are dense.

But the ancients had already proved that those don't include all numbers, e.g. $\sqrt{2}$.

It was shown how reals could be represented as (or modelled by – what's the difference?) bounded convergent infinite sequences of rationals.

Georg Cantor then showed that there are different classes of infinite ordinals and cardinals which could be systematically constructed from the other kinds of numbers.

(Crucial importance of **diagonal** construction: precursor of Gödel, Turing, etc.)

(I won't say anything about complex numbers, e.g. $\sqrt{-1}$.)

Continuity and discontinuity

During the 19th century developments, there were deep interactions between intuitions about continuity (in space and time) and discrete cases and formal (e.g. logical/algebraic) descriptions and reasoning.

For example, integers and ratios (fractions) are discrete entities.

Yet the notion of a continuum appeared to arise out of the way integers and ratios were understood.

Moreover, there were ways of thinking about smooth, continuous changes that imposed a segmentation into distinct sub-cases: for example, the portions of a 2-D curve on either side of a maximum or minimum, or the portions separated by a change of direction of curvature (point of inflection).

> If a curve C has a minimal point P1 (such that all portions of the curve in arbitrarily small neighbourhoods around P1 are higher than P1, and a maximal point P2 (with a similar definition of maximal), and no minimal or maximal points between P1 and P2, then the curve must be monotonically increasing between P1 and P2, and there will be at least one point of maximal steepness of slope between P1 and P2.

As far as I know neither psychologists, neuroscientists nor AI researchers have noticed the variety and importance of such intuitions, and there are no working models of how mathematical discoveries are made about them.

I suspect that nothing currently known by neuroscientists can explain how the mechanisms involved in thinking about proofs in Euclidean geometry work.

I few biologists have noticed the precursors of such mathematical capabilities in the competences of other animals – for example their ability to perceive and use affordances.

Boole, Peano, Frege, Russell, Hilbert

Later developments allowed the intuitive ideas to be expressed in increasingly formal ways and this made it possible to express with great precision ideas about reducing one area of mathematical knowledge to another.

In parallel Frege, Russell and Whitehead (and possibly Peirce?) thought about the possibility of reducing some or all of mathematics to logic, using the recently developed more formal and precise notions of logic.

Russell discovered Frege's work just as the latter was being finished, and found that his logical system allowed Russell's paradox to be derived

> If R is the set of all sets that do not contain themselves then if R contains itself it does not contain itself and if it does not then it does.

That led to a spate of activity attempting to develop a new framework powerful enough to support all of mathematics without entailing a contradiction.

Logic and AI

By the time AI (badly named alas) was founded in 1956 a lot of work had been done developing various formalisations of logic and also algorithmic ideas for checking properties of formulae and proofs.

John McCarthy, AI founder, and inventor of the label "artificial intelligence" had a vision of using logic not just as the foundation of mathematics, but as the foundation of intelligence, at least in its artificial forms.

There have been various attempts to promote alternatives to logicist AI, including use of neural nets, evolutionary computations, rule-based computation not restricted by logical principles, use of statistics and probabilistic reasoning instead of standard logic, and embodied computation where the control mechanisms and physical form and constraints of a robot both contribute to successful actions.

One of the oldest examples is use of a compliant wrist to reduce the need for precise control of a hand.

Many demonstrations of principle and working applications of AI have been produced using these ideas.

But despite many impressive specialised applications, the intelligence of animals interacting with a structured, changing world has not been matched, except for impressive insect-like competences, like those demonstrated by the BigDog robot of Boston Dynamics.

`http://www.bostondynamics.com/robot_bigdog.html`

As far as I can tell, BigDog has no idea what it has done, what has not done what it could have done why it did what it die, what the possible options are that it would face it if performed some action, etc.

But it is very good at maintaining its balance, and gait, and forward motion on a wide variety of terrains.

Annette Karmiloff-Smith's work

So far I have discussed ways in which mathematical reasoning can lead to new knowledge, with the implication that the reasoner (a human mathematician) knows what has been achieved.

But it is possible for an individual to acquire a competence without recognizing or understanding the change.

Examples are children who develop a grasp of the syntax of their language but do not know what they have learnt.

This seems to be an example of what was referred to by Karmiloff-Smith as "Representational redescription" (RR) in Beyond Modularity (Karmiloff-Smith, 1992)

A draft incomplete review of that book is
`http://tinyurl.com/BhamCog/misc/beyond-modularity.html`

She distinguishes different sorts of RR, including

- transforming a competence into a more powerful generative form (like the transition from pattern-based language use to use of syntax)
- acquiring meta-knowledge about the competence, i.e. being able to think about what can and cannot be done
- being able to communicate features of a competence to others who have not yet acquired ti
- being able to perceive the scope and limitations of another's competence

and probably others I've forgotten for now...

Representational re-description and mathematics

It seems likely that humans are not born able to perceive the structures possibilities and invariants involved in the theorems discussed above (e.g. the triangle theorems).

I suspect that that ability depends on first acquiring empirical expertise by exploring the world (compare learning how to talk) and then treating that expertise as itself something to be examined and reorganised by a meta-cognitive mechanism.

That mechanism needs to be able to discern useful patterns that allow new things to be discovered by deriving them from what has previously been learnt, instead of requiring everything to be acquired and tested empirically

> Compare working out how to express a new complex thought using grammatical knowledge, or working out how to create a more stable structure made of bricks.

Conjecture:
Something similar happens in biological evolution
we don't yet know much about the details.

But if we look carefully at pressures to improve competences in perceiving, manipulating, thinking about, forming generalisations about
We may find a succession of evolutionary transitions that lead to human competences.

The difference between online and offline intelligence is crucial, but usually ignored by enactivists, embodiment enthusiasts, and dynamical systems enthusiasts.

Compare Kenneth Craik (Craik, 1943)

Related online stuff

NOTE:

A collection of related talks about meta-morphogenesis which have been or will be added to my "talks" directory:

http://tinyurl.com/BhamCog/talks/

There are several closely related (and still growing) online discussion papers (html) in my 'Misc' directory, e.g.:

- http://tinyurl.com/BhamCog/misc/meta-morphogenesis.html
- http://tinyurl.com/BhamCog/misc/toddler-theorems.html
- http://tinyurl.com/BhamCog/misc/beyond-modularity.html
- http://tinyurl.com/CogMisc/evolution-info-transitions.html

All the above are "work in progess", especially the last web site in transitions in biological information processing, in evolution, development, learning, social systems, ecosystems, ...

They have different, but overlapping, contents.

My paper for the workshop proceedings had to be truncated to fit the space available.

The expanded version (still under development) is online here

http://www.cs.bham.ac.uk/research/projects/cogaff/12.html#1203

Comments on and criticisms of any of this are welcome.

It's a large long-term project.

References

Asperti, A. (2012). Proof, message and certificate. In *Aisc/mkm/calculemus* (p. 17-31). Available from http://dx.doi.org/10.1007/978-3-642-31374-5_2

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

Craik, K. (1943). *The nature of explanation*. London, New York: Cambridge University Press.

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

Jamnik, M., Bundy, A., & Green, I. (1999). On automating diagrammatic proofs of arithmetic arguments. *Journal of Logic, Language and Information*, *8*(3), 297–321.

Kant, I. (1781). *Critique of pure reason*. London: Macmillan. (Translated (1929) by Norman Kemp Smith)

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.

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

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. (http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html)

Sloman, A. (1965). 'Necessary', 'A Priori' and 'Analytic'. *Analysis*, *26*(1), 12–16. Available from Nowonlinehttp://www.cs.bham.ac.uk/research/projects/cogaff/07.html#701

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. (http://www.cs.bham.ac.uk/research/cogaff/04.html#200407)

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. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43)

Sloman, A. (1996). Actual possibilities. In L. Aiello & S. Shapiro (Eds.), *Principles of knowledge representation and reasoning: Proc. 5th int. conf. (KR '96)* (pp. 627–638). Boston, MA: Morgan Kaufmann Publishers. Available from http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15

Sloman, A. (2006, May). *Requirements for a Fully Deliberative Architecture (Or component of an architecture)* (Research Note No. COSY-DP-0604). Birmingham, UK: School of Computer Science, University of Birmingham. Available from http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604

Sloman, A. (2007, Nov). *Predicting Affordance Changes (Alternatives ways to deal with uncertainty)* (Tech. Rep. No. COSY-DP-0702). Birmingham, UK: School of Computer Science, University of Birmingham. Available from http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0702 (Unpublished discussion paper (HTML))

Sloman, A. (2009). Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. In B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, & K. Doya (Eds.), *Creating Brain-like Intelligence* (pp. 248–277). Berlin: Springer-Verlag. Available from http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804

Sloman, A. (2010). If Learning Maths Requires a Teacher, Where did the First Teachers Come From? In A. Pease, M. Guhe, & A. Smaill (Eds.), *Proc. Int. Symp. on Mathematical Practice and Cognition, AISB 2010 Convention* (pp. 30–39). De Montfort University, Leicester. Available from http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1001