School of Computer Science University of Birmingham UK

Hidden Depths of Triangle Qualia

(Especially their areas.)

Theorems About Triangles, and Implications for Biological Evolution and AI The Median Stretch, Side Stretch, and Triangle Area Theorems

Old and new proofs.

NOTE:

Some of the contents of this document, and autobiographical background, are also presented in a video interview with Adam Ford in Oxford, December 2012, available online here: http://www.youtube.com/watch?v=iuH8dC7Snno.

Thanks to help from Dylan Holmes a textual transcript of the interview is now available at http://www.cs.bham.ac.uk/research/projects/cogaff/misc/transcript-agi-interview.html
(Some parts of an earlier version of this document have been moved to the discussion of the triangle sum theorem -- especially Mary Pardoe's proof -- now in: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html)

JUMP TO LIST OF CONTENTS

This file is http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html
A PDF version will be automatically generated from time to time: http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.pdf

---

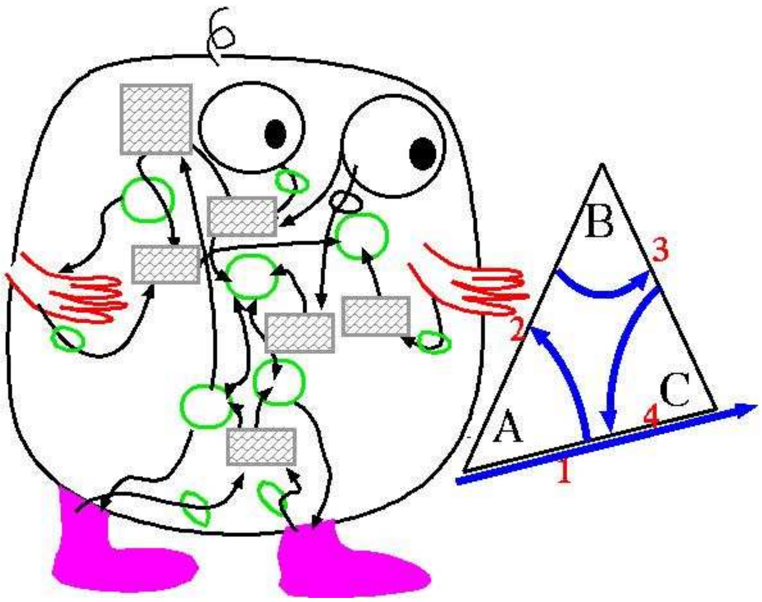

When will the first baby robot grow up to be a mathematician?

CONTENTS

A Preface was added: 7 May 2013, removed 12 May 2013, now a separate document: Biology, Mathematics, Philosophy, and Evolution of Information Processing
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html

---

BACK TO CONTENTS

Introduction: ways to change, or not change, features of a triangle Aspects of mathematical consciousness of space

This document may appear to some to be a mathematics tutorial, introducing ways of doing Euclidean geometry. It may have that function, but my main aim is to draw attention to products of biological evolution that must have existed before Euclidean geometry was developed and organised in Euclid's Elements over two thousand years ago.

An online English version of Euclid's Elements is here:
  http://aleph0.clarku.edu/~djoyce/java/elements/elements.html

That was, arguably, the most important, and most influential, book ever written, ignoring highly influential books with mythical or false contents. Unfortunately, this seems to have dropped out of modern education with very sad results.

I'll give a few examples of apparently very simple, but actually very complex, human spatial reasoning capabilities concerned with perception of triangles. I think these competences are deeply connected with the abilities of human toddlers and other animals to perceive what James Gibson called "affordances", though I don't think he ever understood the full generality, and depth, of those animal competences. I suspect we shall not understand the functions and mechanisms of perception of affordances until we know a lot more about the evolutionary transitions that produced different kinds of affordance-related competences in evolutionary ancestors of humans -- one of the aims of the Meta-Morphogenesis project Sloman(2013).

A core aspect of intelligence of many animals is the ability to perceive what is **possible** in a situation while certain features of the environment remain unchanged. For example it is possible for you to walk through a gap in a wall as long as the width of the gap does not shrink to less than the narrowest width across your body, left to right or front to back. Another core aspect is perception of what is **impossible**/

Perceiving possibilities involves acquiring information about structures and processes that do not exist but could have existed, and might exist in future -- and grasping some of the **constraints** on those possibilities, and also some of the **consequences** if they are realised.

(This is not to be confused with discovering probabilities: possibilities are obviously more basic. The differences between learning about impossibilities, i.e. constraints on what's possible and learning about probabilities seem to have been ignored by most researchers studying probabilistic learning mechanisms, e.g. Bayesian mechanisms. "Deep learning" that draw conclusions from statistical evidence cannot provide information that something is ***impossible****. However, only one instance of a type of configuration or process is required to prove that it is* ***possible****.)*

The examples I'll present look very simple but have hidden depths, as a result of which there is, as far as I know, nothing in AI that is even close to modelling those animal competences, and nothing in neuroscience that I know of that addresses the problem of explaining how such competences could be implemented in brains. How could a neural net encode the information that something (e.g. separation of interlocking rings made of impermeable material without temporarily making a gap in one of the rings, is impossible.)

(I am not claiming that computer-based machines **cannot** model these discovery process, as Roger Penrose sometimes does, only that the current ways of thinking in AI, Computer science, Neuroscience, Cognitive Science, Philosophy of mind and Philosophy of mathematics, need to be extended to allow for perception of possibilities, impossibilities and necessary connections.

I'll be happy to be informed of working models, or even outline designs, for implementing such abilities in robots.

For reasons that should become clearer below, this could be dubbed the problem of accounting for **"mathematical qualia"**, or "contents of mathematical/geometric consciousness" -- their evolution, their cognitive functions, and the mechanisms that implement them.

I have some ideas about the layers of meta-cognitive, and meta-meta-cognitive mechanisms that are involved in these processes, which I think are related to Annette Karmiloff-Smith's ideas about "Representational Redescription", (1992) but I shall not expand on those ideas here: the purpose of this document is to present the problem.

For more on this see the Meta-Morphogenesis project: Sloman(2013), and these examples of perception of impossibilities:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html

An Evolutionary Conjecture

My aim here is to provide examples supporting the following conjecture:

The discoveries organised and presented in Euclid's Elements depended to a considerable extent on products of biological evolution that humans share with several other species of animals that can perceive, understand, reason about, construct, and make use of, structures and processes in the environment -- competences that are also present in pre-verbal humans, e.g. toddlers.

Human toddlers, and some other animals, seem to be able to make such discoveries, but they lack the meta-cognitive competences that enable older humans to inspect and reason about those competences, and the discoveries they give rise to.

I suspect that important subsets of those competences evolved independently in several evolutionary lineages -- including some nest-building birds, elephants, and primates -- because they all inhabit a 3-D environment in which they are able to perceive, understand, produce, maintain, or prevent various kinds of spatial structures and processes. Some of those competences are also present in very young, even pre-verbal, children. But the competences have largely been ignored, or misunderstood, by researchers in developmental psychology, animal cognition, philosophy of mathematics, and more recently AI and robotics. Thinkers who have noticed the gaps sometimes argue that computer-based systems will always have such gaps (e.g. Roger Penrose). That is not my aim, though there is an open question about how research motives should be related to particular sorts of experiments.

NOTE
Research on "tool-use" in young children and other animals often has misguided motivations based on some ill-defined notion of "tool" and the assumption that there is a well defined sub-category of actions that are examples of tool-use.

The label "Tool-use" should perhaps be replaced with "matter manipulation", a label well suited for a wide range of examples. I suspect research on tool use in young children and other animals tends to have an arbitrary focus, and should be replaced by research on "matter-manipulation" including uses of matter to manipulate matter.

But that’s really a topic for another occasion.

Very often these spatial reasoning competences are confused with very different competences, such as abilities to learn empirical generalisations from experience, and to reason probabilistically. In contrast, this discussion is concerned with abilities to discover what is possible, and constraints on possibilities, i.e. necessities.

(These abilities were also noticed by Immanuel Kant, who, I suspect, would have been actively attempting to use Artificial Intelligence modelling techniques to do philosophy, had he been alive now.)

In young humans the mathematical competences discussed here normally become evident in the context of formal education, and as a result it is sometimes suggested, mistakenly, that social processes not only play a role in communicating the competences, or the results of using them, but also determine which forms of reasoning are valid -- a muddle I’ll ignore here, apart from commenting that early forms of these competences seem to be evident in pre-school children and other animals, though experimental tests are often inconclusive: **we need a new deep theory more than we need new empirical data.**

The capabilities illustrated here are, to the best of my knowledge, not yet replicated in any AI system, though some machines (e.g. some graphics engines used in computer games), may appear to have superficially similar capabilities if their limitations (discussed below) are not exposed.

Two disclaimers concerning impossibility/necessity

1. Local vs global impossibility

(Added: 9 Apr 2017)
The examples presented here concern local, or relative, impossibility, not the sorts of **global** impossibility found in pictures by Reutersvard, Penrose, Escher and others, discussed here:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html
Some (Possibly) new considerations regarding impossible objects.

That discusses depictions (and descriptions) of various types of spatial configuration that could not possibly exist because instances would violate some mathematical constraint, such as that *further than* is *transitive* and *irreflexive*.

In contrast the examples here all concern structures and processes that are possible, but which cannot be altered in a way that produces a change satisfying some description. E.g. if X increases then it is impossible for Y to decrease. However, those impossible changes are not ones for which compelling depictions are presented, unlike the impossible objects that have compelling depictions.

2. What CURRENT computer systems cannot do

I am not claiming that computers **cannot** do things done by humans and other animals, merely that novel forms of representation and reasoning, and possibly new information-processing architectures, will be required -- developing a claim I first made in Sloman(1971), though I did not then expect it would take so long to replicate these animal capabilities. That is partly because I did not then understand the full implications of the claims, especially the connection with some of J.J. Gibson's ideas about the functions of perception in animals discussed below, and the distinction between **online intelligence** and **offline intelligence** also discussed below, which challenges some claims made recently about "embodied cognition" and "enactivism", claims that I regard as deeply confused, because they focus on only a subset of competences associated with being embodied and inhabiting space and time.

It is possible that computers will turn out to have a limitation that isn't present in brains but so far there is no evidence that that is the case. An investigation of that possibility is a major feature of the Meta-Morphogenesis project, of which this document is a part Sloman(2013).

> NOTE: The ideas presented here overlap somewhat with ideas of Jean Mandler on early conceptual development in children and her use of the notion of an "image schema" representation, though she seems not to have noticed the need to account for competences shared with other animals. Studying humans, and trying to model or replicate their competences, while ignoring other species, and the precocial-altricial spectrum in animal development, can lead to serious misconceptions.
>
> (I am grateful to Frank Guerin for reminding me of Mandler's work, accessible at http://www.cogsci.ucsd.edu/~jean/)
>
> Another colleague recently drew my attention to this paper:
> http://psych.stanford.edu/~jlm/pdfs/Shepard08CogSciStepToRationality.pdf
> Roger N. Shepard,
> The Step to Rationality: The Efficacy of Thought Experiments in Science, Ethics, and Free Will,
> In *Cognitive Science,* Vol 32, 2008,

It’s one thing to notice the importance of these concepts and modes of reasoning. Finding a good characterisation and developing a good explanatory model are very different, more difficult, tasks.

[Note added: 3 Jan 2013]
This document is also closely related to my 1962 DPhil Thesis attempting to explain and defend Immanuel Kant’s claim (1781) that mathematical knowledge includes propositions that are necessarily true (i.e. it’s impossible for them to be false) but are not provable using only definitions and logic -- i.e. they are not analytic: they are **synthetic necessary** truths.

The thesis is available online in the form of scanned in PDF files, kindly provided by the university of Oxford library, enhanced in 2016 by re-typed searchable (and less bulky) files:

DPhil Thesis (1962)
Aaron Sloman,
Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth
http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962

The most directly relevant section is Chapter 7 "Kinds of Necessary Truth".

The original scanned thesis is available in the Oxford library here.

[End Note]

The Triangle Sum Theorem

Very many people have learnt (memorised) the triangle sum theorem, which states that the interior angles of **any** triangle (in a plane) add up to half a rotation, i.e. 180 degrees, or a straight line, even if they have never seen or understood a proof of theorem. Many who have been shown a proof cannot remember or reconstruct it. A wonderful proof due to Mary Pardoe is presented in http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html For now, notice that the theorem is not at all obvious if you merely look at an arbitrary triangle, such as Figure T:

Figure T:

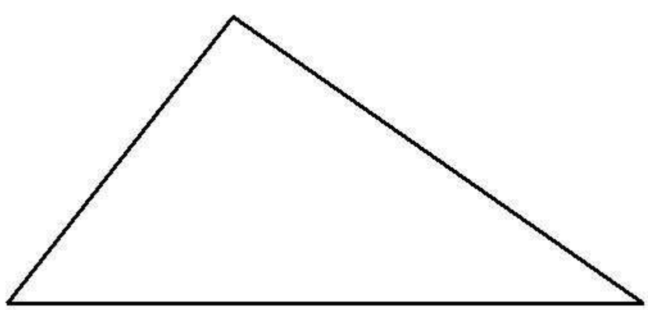

Please stare at that for a while and decide what you can learn about triangles from it. Later you may find that you missed some interesting things that you are capable of noticing.

NOTE:
I don’t know whether the similarity between this exercise and some of the exercises described by Susan Blackmore in her little book **Zen and the Art of Consciousness,** discussed here, is

spurious or reflects some deep connection.

Some ways of thinking about triangles and what can be done with them, including ways of proving the triangle sum theorem, will be presented later. Before that, I'll introduce some simpler theorems concerning ways of deforming a triangle, and considering whether and how the enclosed area must change when the triangle is deformed.

NB: Note that that's "must change", not "will change", nor "will change with a high probability". These mathematical discoveries are about what **must** be the case. Sometimes researchers who don't understand this regard mathematical knowledge as a limiting case of empirical knowledge, with a probability of 1.0. Mathematical necessity has nothing to do with probabilities, but everything to do with constraints on possibilities, as I hope will be illustrated below.

As a warm-up for some of the examples to come later, consider a triangle and what happens to its shape if you keep one side fixed, and move the opposite vertex.

Figure Para

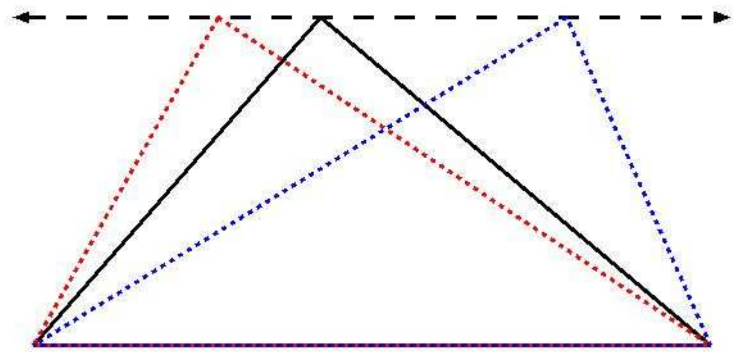

What happens to the size of the top angle of a triangle if the vertex at the top is moved in either direction on a line parallel to the opposite side of the triangle, in this case a line parallel to the base of the triangle?

I suspect that if you consider such changes of the position of the vertex you will see that the angle at the top must change as the vertex slides, but whether the size of the angle will increase or decrease depends on the position at which it starts moving and in which direction it will move.

Is there a location at which the angle will be at its largest size?

An easier problem is presented below. Suppose you move the top vertex of a triangle not as shown in Figure Para namely parallel to its base, but instead move it further from the base along a line that goes through the base and the vertex. Two different examples of this are shown in Figure Move-up, below. What can you say about the size of the top angle as the top vertex moves further from the base along a straight line? Two different examples of such motion are shown in the figure.

Figure Move-up

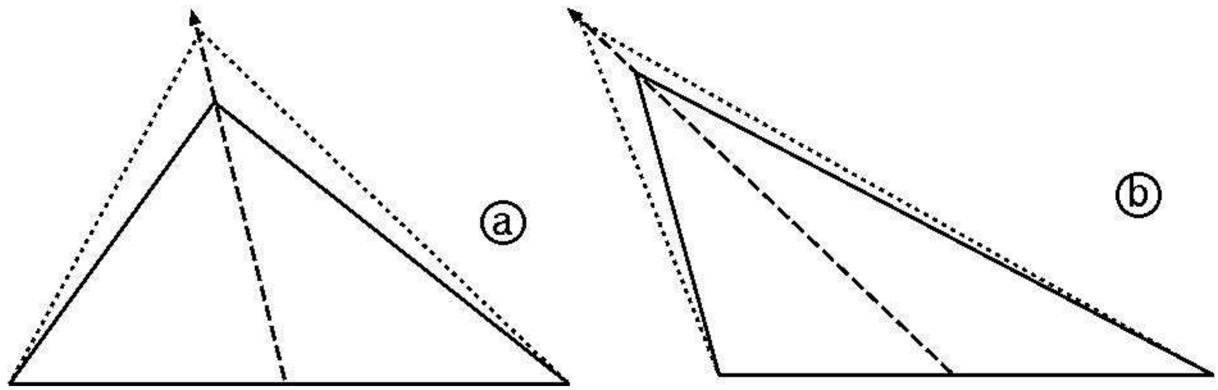

Can you give a reason for your answer? Is there something about the diagram that makes it obvious that moving the top vertex away from the opposite side will decrease the size of angle at the top?

As far as I know this question was not asked by Euclid or discussed in the Elements, but it is closely related to some of the conditions under which two triangles are thought to be necessarily congruent, so that either can be placed over the other without leaving any portion of the other visible.

I offer the problems posed by Figure-Para and Figure Move-up, because most people will find the answers obvious (especially the second for which the size of the upper vertex must decrease as it moves away from the base.)

I shall try later, in another document, with a link added here, to give an analysis of requirements for a machine to make such discoveries and to be able to conclude that violations to the generalisation are *impossible*. In Immanuel Kant's terms: what has been discovered is (a) non-empirical, (b) necessarily true, and (c) is not a logical consequence of definitions of terms. See Sloman(1962).

Why focus on the human ability to notice and prove some invariant property of triangles? Because it draws attention to abilities to perceive and understand things that are closely related to what James Gibson called "affordances" in the environment, namely: animals can obtain information about possibilities for action and constraints on action that allow actions to be selected and controlled. An example might be detecting that a gap in a wall is too narrow to walk through normally, but not too narrow if you rotate your torso through a right angle and then walk (or sidle) sideways through the gap. **NB:** You can notice the possibility, and think about it, without making use of it. Use of offline intelligence is neither a matter of performing actions at the time or in the immediate future, nor making predictions. The ability to discover such a possibility is not always tied to be able to make use of it in the near future.

Video 6 here illustrates an 19 month old toddler's grasp of affordances related to a broom, railings and walls:
http://www.cs.bham.ac.uk/research/projects/cogaff/movies/vid

The video, shows the child manipulating a broom, and includes a variety of actions in which the child seems to understand the constraints on motion of the broom and performs appropriate actions, including moving it so as to escape the restrictions on motion that exist when the broom handle is between upright rails, moving the broom backwards away from a skirting board in order to

be able to rotate it so that it can be pushed down the corridor, and changing the orientation of the vertical plane containing the broom so that by the time it reaches the doorway on the right at the end of the corridor the broom is ready to be pushed through the doorway.
NB: I am not claiming that the child understands what he is doing, or proves theorems.

Video 8 shows a younger child (17.5 months) apparently formulating and testing conjectures in 3D topology regarding possible routes in 3D space between different configurations of a pencil and sheet of card with a hole in it. I suspect no current theory in neuroscience, developmental psychology or linguistics, explains how her (pre-verbal) brain is able to represent the contents of her percepts, intentions, hypotheses and plans. Similar questions arise for mental states and processes in other intelligent animals that do not use human languages, e.g. elephants, orangutans, squirrels, weaver birds, and others.

BACK TO CONTENTS

Doing without a global length metric

A feature of such abilities, whose importance should become clearer later, is that they do not depend on the ability to produce or use accurate measurements, using a global scale of length, or area. They do depend on the ability to detect and use ordering information, such as the information that your side-to-side width is greater than the width of a gap you wish to go through, and that your front-to-back width is less than the width of the gap, and your ability to grasp the possibility of rotating and then moving sideways instead of always moving forwards. A **partial** order suffices: you don’t need to be able to determine for all gaps viewed at a distance whether your side-to-side or front-to-back distance exceeds the gap.

Using the ordering information (when available), you can infer that although forward motion through the gap is impossible, sideways motion through it is possible. It is important that your understanding is not limited to exactly this spatial configuration (this precise gap width, this precise starting location, this precise colour of shoe, this kind of floor material on which you are standing), since you can abstract away from those details to form a generic understanding of a class of situations in which a problem can arise and can be solved. The key features of the situation are relational, e.g. the gap is narrower than one of your dimensions but greater than another of your dimensions. It does not require absolute measurements. If you learn this abstraction as a child confronted with a particular size gap you can still use what you have learnt as an adult confronted with a larger gap, that the child could have gone through by walking forward.

BACK TO CONTENTS

Offline vs online intelligence

That sort of abstraction to a general schema that can be instantiated in different ways is at the heart of mathematical reasoning (often confused with use of metaphor). It is also important for **offline** intelligence (reasoning and about what can be done, and planning) as opposed to **online** intelligence (used by servo-controlled reactive and homeostatic, systems) a distinction discussed further in another document.

Observation of actions of many different animals, for instance nest building birds, squirrels, hunting mammals, orangutans moving through foliage, and young pre-verbal children, indicates to the educated observer (especially observers with experience of the problems of designing intelligent robots) many animal capabilities apparently based on abilities to perceive, understand and use affordances, some of them more complex than any discussed by Gibson, including "epistemic affordances" concerned with possible ways of gaining information, illustrated here. For example, If you are in a corridor outside a room with an open door, and you move in a straight line towards the centre of the doorway, you will see more of the room, and will therefore have access to more information, an epistemic affordance. (This illustrates the connection between theorems in Euclidean geometry and visual affordances that are usable by humans, other animals, and future robots.)

I suspect, but will not argue here, that the human ability to make mathematical discoveries thousands of years ago, that were eventually gathered into a system and published as Euclid's Elements, depended on the same capability to discover affordances, enhanced by additional meta-cognitive abilities to think about the discoveries, communicate them to others, argue about them, and point out and rectify errors.

NOTE:
(These social processes are important but sometimes misconstrued, e.g. by conventionalist philosophers of mathematics. They will not be discussed here.)

The information-processing (thinking) required in offline intelligence is sometimes too complex to be done entirely within the thinker, and this may have led to the use of external information structures, such as diagrams in sand or clay or other materials, to facilitate thinking and reasoning about the more complex affordances, just as modern mathematicians use blackboards, paper and other external thinking aids, as do engineers, designers, and artists. (As discussed in (Sloman, 1971).)

The roles of external representational media in discovering re-usable generalisations is different from their meta-cognitive role in reasoning about the status of those generalisations, e.g. proving that they are **theorems**. That difference is illustrated but not explained or modelled in this document.

In this document I have chosen some very simple, somewhat artificial, cases, simply to illustrate some of the properties of offline thinking competences, in particular, how they differ from the ability of modern computer simulation engines that can be given an initial configuration from which they compute in great detail, with great precision what will happen thereafter. That simulation ability is very different from the ability to think about a collection of possible trajectories, features they have in common and ways in which they differ, a requirement for the ability to create multi-stage plans. (I am not sure Kenneth Craik understood this difference when he proposed that intelligent animals could use internal models to predict consequences of possible actions, in (Craik, 1943).)

There are attempts to give machines this more general ability to learn about and use affordances, by allowing them to learn and use probability distributions, but I shall try to explain below why that is a very different capability, which lacks the richness and power of the abilities discussed here, though it is sometimes useful. In particular, the probability-based mechanisms lack the ability required to do mathematics and make mathematical discoveries of the kinds illustrated below and in other documents, including "toddler theorems" of the sorts pre-verbal children seem able to discover and use, though there are many individual differences between children: not all can discover the same theorems, nor do they make discoveries in the same order.

One of several motivations for this work is to draw attention to some of what needs to be explained about biological evolution, for the capabilities discussed here all depend on evolved competences -- though some may also be implemented in future machines.

BACK TO CONTENTS

Flaws in enactivism and theories of embodied cognition

Another motivation is to demonstrate that some researchers in AI/Robotics, cognitive science and philosophy have been seriously misled by the recent emphases on embodiment, and enactivist theories of mind, which are mostly concerned with "online intelligence" and ignore the varieties of "offline intelligence" that become increasingly important as organisms grow larger with more complex and varied needs. Some varieties of offline intelligence (sometimes referred to as deliberative intelligence) required for geometrical reasoning are discussed below. A broader discussion is Sloman (2006)

The small successes of the embodied/enactivist approaches (which are, at best, barely adequate to explain competences of some insects) have diverted attention from the huge and important gaps in our understanding of animal cognition and the implications for understanding human cognition and producing robots with human-like intelligence. Although the online intelligence displayed by the BigDog robot made by Boston Dynamics [REF] is very impressive, it remains insect-like, though perhaps not all insects are restricted to "online" intelligence, which involves reacting to the environment under the control of sensorimotor feedback loops, in contrast with "offline" intelligence, which involves being able to consider, reason about and make use of possibilities (not to be confused with probabilities) some of which are used and some avoided. Related points were made two decades ago by David Kirsh, (though he mistakenly suggests that tying shoelaces is a non-cerebral competence, possibly because it can become one through training, as can many other competences initially based on reasoning about possibilities).

I’ll now return to mathematical reasoning about triangles, hoping that readers will see the connection between that and the ability to use offline intelligence to reason about affordances. In what follows, I’ll refer to what **you** can see, on the assumption that all the people reading this document are likely to share some very basic cognitive competences of the sort that preceded the development of Euclidean geometry. If it turns out that you cannot see some of the things I describe please email me with a summary of the problem.

BACK TO CONTENTS

On seeing triangles (again)

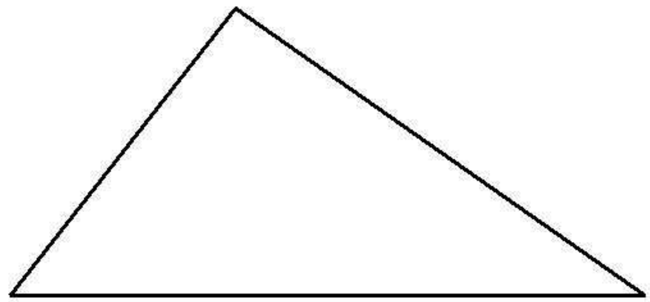

Figure T (repeated):

When you look at a diagram like Figure T, above, you are able to think of it as representing a whole class of triangles with different shapes, sizes, orientations, colours, etc., and you can also think about processes that change some aspect of the triangle, such as its shape, size, orientation, or area, in a way that is not restricted to that particular triangle.

This is very different from the suggestion (by Bishop Berkeley?) that thinking about all triangles requires use of a blurred image produced by superimposing images of many or all possible shapes and sizes. That would be useless for our purposes.

How you think about and reason about possible changes in a spatial configuration is a deep question, relevant to understanding human and animal cognition -- e.g. perception of and reasoning about spatial affordances, and also relevant to the task of designing future intelligent machines. Several examples will be presented and discussed below.

As far as I know, there is no current AI or robotic system that can perform these tasks, although many can do something superficially similar, but much less powerful, namely answer questions about, or make predictions about, a very specific process, starting from precisely specified initial conditions. That is not the same as having the ability to reason about an infinite variety of cases. Machines can now do that using equivalent algebraic problems, but they don’t understand the equivalence between the algebraic and the geometric problems, discovered by Descartes.

> **Note (added 9 Nov 2016)**
> Similar questions can be asked about what happens when you look at an abstract logical formula and think about all possible ways of replacing the variables with constants. It is not often noticed that logical reasoning using written formulae depends on the ability to perceive and reason about spatial structures, and possible ways of combining or modifying them.
>
> Computers can be programmed to do something similar, but what they do and how they do it is different in interesting ways. They do not simulate or replicate human logical reasoning when they perform logical inferences.

The perception of possible changes in the environment, and constraints on such changes, is an important biological competence, identified by James Gibson as perception of "**affordances**". However, I think he noticed and understood only a small subset of types of affordance. His ideas are presented and generalised in a presentation on his ideas (and Marr’s ideas) mentioned above.

I shall present several examples of your ability to perceive and reason about possibilities for change, and constraints on those possibilities inherent in a spatial configuration, extending the discussion in my 1996 "Actual Possibilities" paper.

In particular, we need to discuss your ability to:

a) perceive a shape,
b) notice the possibility of a certain constrained transformation of that shape
c) discover and prove a consequence of that constraint.

Such mathematical competences seem to be closely related to much more wide-spread animal competences involving perception of possibilities for change, including possibilities for action in the environment; and reasoning about consequences of realising those possibilities. The mathematical competences build on these older, more primitive, competences, which seem largely to have gone unnoticed by researchers in human and robot cognition. I have tried to draw attention to examples that can be observed in young children in a discussion of "toddler theorems".

Several proofs of simple theorems will be presented below, making use of your ability to perceive and reason about possible changes in spatial configurations. I'll start with some deceptively simple examples relating to the area enclosed by a triangle.

A developmental neuroscience researcher whose work is closely related to this is Annette Karmiloff-Smith, especially the ideas about "Representational Redescription" in her book "Beyond Modularity" (1992).

BACK TO CONTENTS

The "Median Stretch Theorem"

The first theorem concerns the consequences of moving one vertex of a triangle along a median, while the other two vertices do not move. I shall start by assuming that the concept of the **area** enclosed by a set of lines is understood, and that at least in some cases we can tell which of two areas is larger. Later, I'll return to hidden complexity in the concept of area.

A median of a triangle is a straight line between the midpoint of one side of the triangle to the opposite vertex (corner). The dashed arrows in Figure M (a) and (b) lie on medians of the triangles composed of solid lines. The dashed arrows in triangles (a) and (b) have both been extended beyond the median, which terminates at the vertex. The dotted lines indicate the new locations that would be produced for the sides of the triangle if the vertex were moved out, as shown.

Figure M:

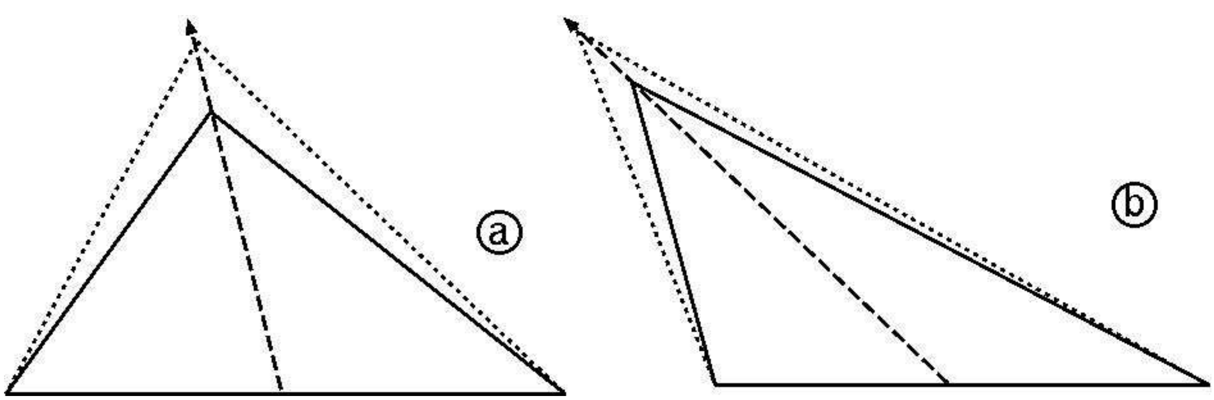

Consider what happens if we draw a median in a triangle, namely a line from the midpoint of one side through the opposite vertex, and then move that vertex along the extension of the median, as shown in figure M(a). You should find it very obvious that moving the vertex in one direction along the median increases the area of the triangle, and movement in the other direction decreases the area. Why?

We can formulate the "Median stretch theorem" (MST) in two parts:

> **(MST-out)** IF a vertex of a triangle is moved along a median **away from** the opposite side, THEN the area of the triangle **increases**.
>
> **(MST-in)** IF a vertex of a triangle is moved along a median **towards** the opposite side, THEN the area of the triangle **decreases**.

As figure M(b) shows, it makes no difference if the vertex is not perpendicularly above the opposite side: the diagrammatic proof displays an invariant that is not sensitive to alteration of the initial shape of the triangle, e.g. changing the slant of the median, and changing the initial position of the vertex in relation to the opposite side makes no difference to the truth of theorem. Why?

A problem to think about: How can you be sure that there is no counter-example to the theorem, e.g. that stretching or rotating the triangle, or making it a different colour, or painting it on a different material, or transporting it to Mars, will not make any difference to the truth of (MST-out) or (MST-in)?

> **NOTE:** As far as I know the median stretch theorem has never been stated previously, though I suspect it has been used many times as an "obvious" truth in many contexts, both mathematical and non-mathematical. If you know of any statement or discussion of the theorem, please let me know.
>
> **Note added 24 Feb 2013:** Readers may find it obvious that the median stretch theorem is a special case of a more general stretch theorem that can be formulated by relaxing one of the constraints on the lines in the diagram. Figuring out the generalisation is left as an exercise for the reader. (Feel free to email me about this.) Compare (Lakatos, 1976).
>
> **Added 13 Feb 2013** Julian Bradfield pointed out, in conversation, that one way to think about the truth of MST-OUT is to notice that the change of vertex **adds two triangles** to the original triangle. Likewise, in support of MST-IN, moving the vertex inwards **subtracts** two triangles from the original area. This also applies to the MCT (containment) theorems, below. (Below I suggest decomposing the proof into two applications of the Side-Stretch-Theorem (SST), which can also be thought of as involving the addition or subtraction of a triangle.)

BACK TO CONTENTS

NOTE: How can areas be compared?

> The concept of "area" used here may seem intuitive and obvious, but generalising it to figures with arbitrary boundaries is far from obvious and requires the use of sophisticated mathematical reasoning about limits of infinite sequences.
>
> For example, how can you compare the areas of an ellipse and a circle, neither of which completely encloses the other? What are we asking when we ask whether the blue circle or the red ellipse has a larger area in Figure A, below? It is obvious that the black square contains less space than the blue circle, and also contains less space than the red ellipse, simply because all the space in the square is also in side the circle and inside the ellipse. But

what does it mean to ask whether one object contains more space than another if each cannot fit inside the other?

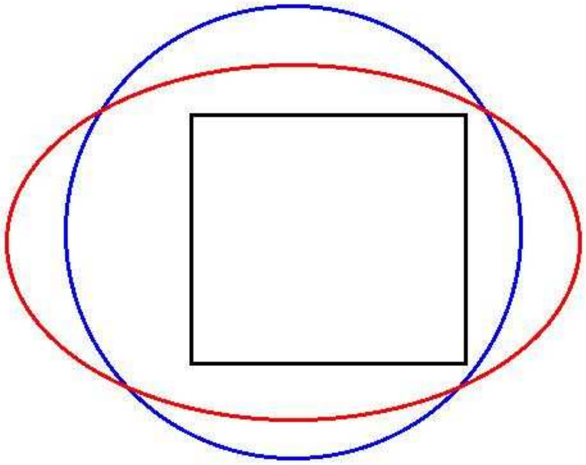

Figure A:

Some teachers try to get young children to think about this sort of question by cutting out figures and weighing them. But that assumes that the concept of weight is understood. In any case we are not asking whether the portion of paper (or screen!) included in the circle weighs more than the portion included in the ellipse. There is a correlation between area and weight (why?) but it is not a reliable correlation. Why not?

The standard mathematical way of defining the area of a region includes imagining ways of dividing up non-rectangular regions into combinations of regions bounded by straight lines (e.g. thin triangles, or small squares), using the sum of many small areas as an approximation to the large area. The smaller the squares the better the approximation, in normal cases. (Why? -- Another area theorem).

For our purposes in considering the triangles in Figure M and Figure S, most of those difficulties can be ignored, since we can, for now, use just the trivial fact that if one region totally encloses another then it has a larger area than the region it encloses, leaving open the question of how to define "area", or what it means to say that area A1 is larger than area A2, when neither encloses the other. Our theorems about stretching (MST above and SST below) only require consideration of area comparisons when one area completely encloses another.

Cautionary note: It is very easy for experimental researchers studying animals or young children to ask whether they do or do not understand areas (or volumes, or lengths of curved lines), and devise tests to check for understanding, without the researchers themselves having anything like a full understanding of these concepts that troubled many great mathematicians for centuries. (I have checked this by talking to some of the researchers, who had not realised that the resources for thinking about areas and volumes in very young children might support only a **partial** ordering of areas.

These problems are usually made explicit only to students doing a degree in mathematics.

Using containment instead of area (Added 27 May 2013)

(Related to Julian Bradfield's observation, above.)
In order to avoid the concept of area we can switch to a more basic theorem about **containment**, which, once again has two parts. After the vertex is moved along the median a new triangle is obtained. Instead of thinking about the **areas** of the old and the new triangles we can ask about whether one is **contained in** the other. We get two new theorems, which are more primitive than the stretch theorems as stated above. We can formulate the "Median Containment Theorem" **(MCT)** in two parts:

> **(MCT-out)** IF a vertex of a triangle is moved along a median **away from** the opposite side, THEN the new triangle obtained **contains** the original triangle.
>
> **(MCT-in)** IF a vertex of a triangle is moved along a median **towards** the opposite side, THEN the new triangle obtained **is contained within** the original triangle.

In fact, it appears that the ability to see that MST is true depends on the ability to see that MCT is true. Later we'll consider examples where MST remains true, but MCT does not, so that a more sophisticated form of reasoning, using the concept of area is required.

The word "includes" could be used instead of "contains". There are probably several more equivalent formulations that do not make use of the concept of a *measurable area of a particular triangle*, but instead use the concept of a *relation between two triangles* which has nothing to do with measurement, or numbers.

Is this geometry or topology? (27 May 2013) A formulation that does not mention length or area, only one line segment **containing** another, and a triangle **containing** another, comes close to using only topological concepts, though the notion of straightness is still used, and that is not a topological concept. There is a subset of Euclidean geometry that is topology enhanced with notions of straightness and planarity, though officially topology (often confused with topography by non-mathematicians) was not started as a branch of mathematics until centuries after Euclid. See http://en.wikipedia.org/wiki/Topology#History

The task of finding a proof without using metrical relations is discussed further below, in connection with the Side Containment Theorem (SCT).

BACK TO CONTENTS

Using the Side Stretch Theorem to prove the Median Stretch Theorem (Modified: 27 May 2013)

In this section, we'll introduce the Side Stretch Theorem (SST) and show how it was implicitly assumed in the proof of the Median Stretch and Median Containment Theorems (MST and MCT), above.

Thinking about why the MST and MCT must always be true requires noticing that each of the two triangles being compared (the original one before the stretch, and the new one after the stretch, using different portions of the extended median) is made of two smaller triangles, one on each side of the median (shown on each side of the dashed line in Figures M (a) and (b), above).

So the process of comparing sizes of the larger triangles before and after the change of vertex can be broken down into two parts, one on each side of the median. If the area of each **part** of the triangle is increased or decreased in the same way when the location of the vertex on the median changes, then the area of the **whole** triangle composed of the two parts must be increased or decreased. Moreover, if we wish not to assume the concept of Area we can simply focus on containment as in the Median Containment Theorem (MCT) above.

If you think about your reasoning about the change in area of each of the two sub-triangles, you may notice the implicit use of another theorem, which could be called "The side stretch theorem" (SST) illustrated in figure S -- later generalised to the Side Containment Theorem (SCT) here

Figure S:

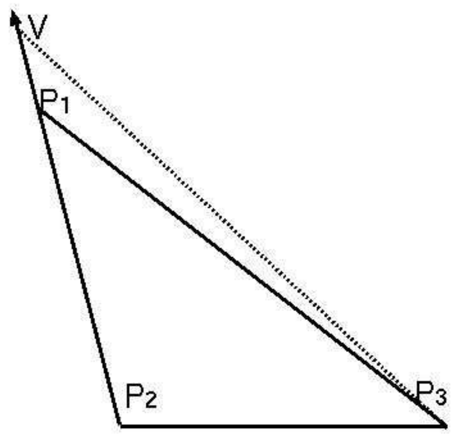

We can formulate the "Side stretch theorem" (SST) in two parts:

(SST-out) IF a vertex of a triangle is moved along an extended side **away** from the interior of the side (as in Figure S) THEN the area of the triangle increases.

(SST-in) IF a vertex of a triangle is moved along a side **towards** the interior of that side, THEN the area of the triangle decreases. (Draw your own figure for this case.)

Comparing Figure S, with Figure M (a) or Figure M (b) should make it clear that when a vertex moves along the median of either of the triangles in Figure M, then there are also two smaller triangles, each of which has one side on the median, and when the vertex of the big triangle moves along the median then the (shared) vertex of each of the smaller triangles moves along the shared side.

Moreover, when the shared vertex in Figure M (a) or (b) moves along the median, both of the smaller triangles either increase or decrease in area simultaneously (in accordance with SST-out or SST-in), from which it follows that their combined area must increase when the vertex moves along the median **away from** the opposite side and decrease when the vertex moves along the median **towards** the opposite side.

The Side Containment Theorem (SCT) It might be fruitful for the reader to pause here and try to formulate and prove a Side Containment Theorem (**SCT**) expressed in terms of which of two triangles contains the other, without assuming any measure of area or length. This can be modelled on the transition presented above, from MST, mentioning stretching and areas, to MCT, mentioning only containment of lines and triangles, not length or area.

For now, I’ll leave open the question whether the Side Stretch Theorem (SST-in/out), or the Side Containment Theorem (SCT) can be derived from something more basic and obvious, requiring biologically simpler, evolutionarily older, forms of information processing.

Note, however, that the notion of the vertex being "moved along" a line "away from" or "towards" another point on the line implicitly makes use of a metrical notion of length, which increases or decreases as the vertex moves. The concept of motion between two locations on a line also implicitly assumes the existence of intermediate locations between those locations.

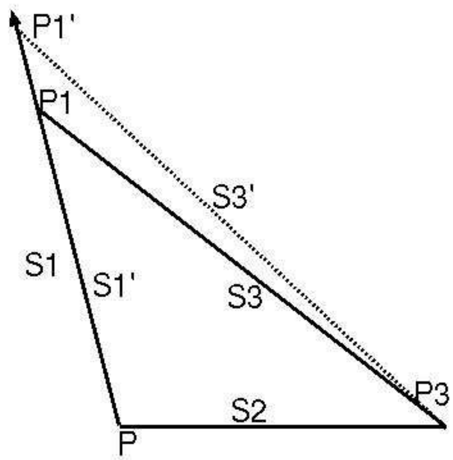

Figure SCT:

It is possible to adopt a different view of this problem by avoiding mention of time or motion, and instead simply referring to two line segments, S1 contained in S1’, with a shared end point P, their other ends being P1 and P1’, as shown in Figure SCT; with another line segment S2 also sharing an end point with S1 and S1’, and two triangles completed with segments S3 between P3 and P1 and segment S3’, between P3 and P1’, as shown in Figure SCT.

In this situation there are two triangles that can be considered. Triangle T which is bounded by segments S1, S2 and S3, with point P1 as vertex; and Triangle T’ which is bounded by Segments S1’, S2 and S3’ and with point P1’ as vertex.

Then theorem SCT states that if segment S1’ contains segment S1, as shown, then triangle T’ will contain triangle T, as should be clear from the diagram. Notice that we have removed reference to motion and time, but considered only containment relations between static objects, two line segments and two triangles containing those segments. By analogy with theorem SCT we can collapse the two theorems MCT-out and MCT-in into one theorem MCT which states that if S1’ contains S1 then the triangle formed from S1’ contains the triangle formed from S1.

Invariants over sets of possibilities The relationship between direction of motion of the vertex V and whether the area increases or decreases, or the changes in containment corresponding to motion of V, can be seen to be **invariant** features of the processes. But it is not clear what information-processing mechanisms make it possible to discover that invariance, or necessity.

We are not discussing probabilities here, only what’s possible or impossible Notice that this is utterly different from the kind of discovery currently made by AI programs that collect large numbers of observations and then seek statistical relationships in the data generated, which is how much robot learning is now done. The kind of learning described here, when done by a human, does not require large amounts of data, nor use of statistics. There are no probabilities involved, only invariant relationships: if a vertex P1 moves along one side, away from the opposite end of that side P2, and the other two vertices, P2 and P3 do not move, then the new triangle must contain the old one. This is not a matter of a high probability, not even 100% probability. It’s about what combinations of states and processes are **impossible**.

The roles of continuity and discontinuity in SST and SCT As a vertex V moves along one side **away** from the other end of that side, P2, the area will increase continuously. However, if the vertex moves in the opposite direction, **towards** the other end, i.e. V moves toward P2, then the change in direction of motion necessarily induces a change in what happens to the area: instead of **increasing**, the area must **decrease**.

As the vertex moves with continuously changing location, velocity, and/or acceleration, there are some unavoidable **discontinuities**: the direction of motion can change, and so can whether the area is increasing or decreasing. But there are more subtle discontinuities, which can be crucial for intelligent agents.

Virtual discontinuities (28 May 2013) A "virtual discontinuity" can occur during continuous motion with fixed velocity and direction. If a vertex V of a triangle starts beyond a position P1 on one of the sides of the triangle and moves back towards the other end of the line, P2, then the location of V, the distance from the other end, and the area of triangle VP2P3 will all change **continuously**.

But, for every position P on the line through which V moves, there will be a **discontinuous** change from V being further than P from the opposite end (P2), to V being nearer than P to the opposite end. Likewise, there will be a discontinuous change from **containing** the triangle with vertex P to **being contained** in that triangle, and the area of the triangle with vertex V will change discontinuously from being greater than the area of the triangle with vertex P to being smaller. Between being greater, or containing, to being smaller, or being contained there is a state of "instantaneous equality" separating the two phases of motion.

This discontinuity is not **intrinsic** to the motion of V, but involves a **relationship** to a particular point P on the line. The same continuous motion can be interpreted as having different virtual discontinuities in relation to different reference points on the route of the change.

If there is an observer who has identified the location P the discontinuity may be noticed by the observer. But there need not be any observer: the discontinuity is there in the space of possible shapes of the triangle as V moves along one side.

There are many cases where understanding mathematical relationships or understanding affordances involves being able to detect such virtual discontinuities based on relational discontinuities (phase changes of a sort). For example, a robot that intends to grasp a cylinder may move its open gripper until the ’virtual’ cylinder projected from its grasping surfaces down to the table contains the physical cylinder. Then it needs to move downwards until the gripping surfaces are below the plane of the top surface of the cylinder, passing through another virtual discontinuity. Then the gripping surfaces can be moved together until they come into contact with the surfaces of the cylinder: a physical, non-virtual, discontinuity. An expert robot, or animal, instead of making the three discrete linear motions could work out (or learn) how to combine them into a smooth curved trajectory that subsumes the three types of discontinuity. But without understanding the requirements to include the virtual discontinuities a learning robot could waste huge amounts of time trying many smooth trajectories that have no hope of achieving a grasp.

Note: After discovering this strategy in relation to use of one hand a robot or animal may be able to use it for the other hand. Moreover, since the structure of the trajectory and the conditions for changing direction are independent of whose hand it is, the same conceptual/perceptual apparatus can be used in perceiving or reasoning about the grasping action of another individual capable of grasping with a hand. It may even generalise to other modes of grasping, e.g. using teeth, if the head can rotate. (The concept of a "mirror neuron" may be be unwittingly based on the assumption that individual neurons can perform such feats or control, or perception.)

BACK TO CONTENTS

Another way of modifying a triangle

Instead of considering what happens when we move the upper vertex in Figure T so that it moves along a median or along a side of the triangle, we can consider possible changes in which the vertex remains at the **same** distance from the opposite side, which would be achieved by moving it along a line **parallel** to the opposite side instead moving it along a median a side of the triangle. (Note: The notion of parallelism includes subtleties that will be ignored for now.)

In Figure Para, below, two new dotted triangles have been added to the triangle formerly shown in Figure T: a new red one and a new blue one, both with vertices on the dashed line, parallel to the base of the original triangle, and both sharing a side (the base) with the original triangle.

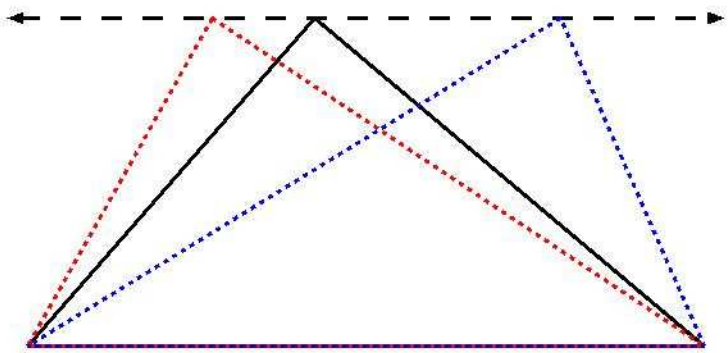

Figure Para:

The figure shows that moving the top vertex of the original triangle along a line **parallel** to the opposite side will definitely not produce a triangle that encloses the original, because, whichever way the vertex is moved on a line parallel to the opposite site (the dashed black line in Figure Para) the change produces a triangle with two new sides, one partly inside the old triangle and the other outside the old triangle. So the new triangle cannot enclose the old one, or be enclosed by it.

We therefore cannot read off from this diagram any answer to the question how sliding a vertex of a triangle along a line parallel to the opposite side affects the area. Proving the theorem that moving a vertex of a triangle in a direction **parallel** to the opposite side does not alter the area is left as an exercise for the reader, though I shall return to it below.

There is a standard proof used to establish a formula for the area of a triangle, which requires consideration of different configurations, as we’ll see below. (The need for case analysis is a common feature of mathematical proof Lakatos 1976).

Exercise for the reader:
Try to formulate a theorem about what happens to the sides of a triangle if a vertex moves along a line that goes through the vertex but does not go through the triangle, like the dashed line in Figure Para, above.

BACK TO CONTENTS

Some observations on the above examples

The examples above show that many humans looking at a triangle are not only able to see and think about the particular triangle displayed, but can also use the perceived triangle to support thinking and reasoning about large, indeed infinitely large, sets of **possible** triangles, related in different ways to the original triangle.

Note:
The concept of an infinitely large set being used here is subtle and complex and (as Immanuel Kant noted) raises deep questions about how it is possible to grasp such a concept. For the purposes of this discussion it will suffice to note that if we are considering a range of cases and have a means of producing a new case different from previously considered cases, then that supports an unbounded collection of cases.

For example, in Figure S, where a vertex of a triangle is moved along an extension of a side of the triangle, between any two positions of the vertex there is at least one additional possible position, and however far along the extended side the vertex has been moved there are always further locations to which it could be moved.

Anyone who is squeamish about referring to infinite sets can, for our purposes, refer to unbounded sets.

Below I'll discuss some implications for meta-cognition in biological information processing.

In some cases the new configurations thought about include additional geometrical features, specifying constraints on the new triangle, for example the constraint that a vertex moves on a median of the original triangle, or on an extension of a median, or on a line parallel to one of the sides. Such constraints, involving lines or circles or other shapes, can be used to limit the possible variants of a starting shape, while still leaving infinitely many different cases to be considered.

However, the infinity of possibilities is reduced to a small subset of cases by making use of common features, or invariants, among the infinity of cases.

For instance the common feature may be a vertex lying on a particular line, such as a median of the original triangle (as in the Median Stretch Theorem (MST) above, or an extended side of the original triangle (as in the Side Stretch Theorem (SST above)). In such situations we can divide the infinity of cases of change of length to two subsets: a change that increases the length and a change that decreases the length, as was done for each of the theorems. Each subset has an invariant that can be inspected by a perceiver or thinker with suitable meta-cognitive capabilities, discussed further below.

There are more complex cases, for example a vertex moving on a line parallel to the opposite side of the triangle, or a vertex moving on a line perpendicular to the opposite side. There are infinitely many perpendiculars to a given line, adding complexity missing in the previous examples. Moreover, in all the figures required for discoveries of the sorts we have been discussing, there is an additional infinity of cases because of possible variations in the original triangle considered, before effects of motion of a vertex are studied.

BACK TO CONTENTS

The role of meta-cognition

This ability to think about infinitely many cases in a finite way seems to depend on the biological **meta-cognitive ability** to notice that members of a set of perceived structures or processes share a common feature that can be described in a meta-language for describing spatial (or more generally perceptual) information structures and processes. An example would be noticing that between **any** two stages in a process there are intermediate stages, and that between any two locations on a line, thicknesses of a line, angles between lines, amounts of curvature, there are **always** intermediate cases, with the implication that there are intermediate cases between the intermediate cases and the intermediate cases never run out.

> **NOTE**: for now we can ignore the difference between a set being dense and being continuous -- a difference that mathematicians did not fully understand until the 19th Century. I shall go on referring loosely to 'continuity' to cover both cases.

This ability to notice that some perceived structure or process is continuous, and therefore infinite, is meta-cognitive insofar as it requires the process of perceiving, or imagining, a structure or process to be monitored by another process which inspects the changing information content of what is being perceived, or imagined, and detects some feature of this process such as continuity, or such as being divisible into discrete cases (e.g. motion away from or towards a line). A more complex meta-cognitive process may notice an invariant of the perceived structure or process, for instance detecting that a particular change necessarily produces another change, such as increasing area, or that it preserves some feature, e.g. preserving area.

> **NOTE:** The transitions in biological information processing required for organisms to have this sort of meta-cognitive competence have largely gone unnoticed. But I suspect they form a very important feature of animal intelligence that later provided part of the basis for further transitions, including development of meta-meta-meta... competences required for human intelligence. (Chappell&Sloman 2007)

These meta-cognitive abilities are superficially related to, but very different from, abilities using statistical pattern recognition techniques to cluster sets of measurements on the basis of co-occurrences. Examples of non-statistical competences include being able to notice that certain differences between cases are irrelevant to some relationship of interest, or being able to notice a way of partitioning a continuous set of cases into two or more non-overlapping sub-sets, possibly with partially indeterminate (fuzzy) boundaries between them. In contrast, many of the statistical techniques require use of large numbers of precise measures in order to detect some pattern in the collection of measures (e.g. an average, or the amount of deviation from the average, or the existence of clusters).

For example, you should find it obvious that the arguments used above based on Figure M and Figure S to prove the Median Stretch and Side Stretch theorems (MST and SST) do not depend on the sizes or shapes of the original triangles. So the argument covers infinitely many different triangular shapes. The features that change if a vertex is moved away from the opposite side along a median or along a side will always change in the same direction, namely, increasing the area.

Noticing an invariant topological or geometrical relationship by abstracting away from details of one particular case is very different from searching for correlations in a large number of particular cases represented in precise detail. For example computation of averages and various other statistics requires availability of many particular, precise, measurements, whereas the discovery process demonstrated above does not require even one precisely measured case. The messy and blurred Figure S-b will do just as well to support the reasoning used in connection with the more precise Figure S, though even that has lines that are not infinitely thin.

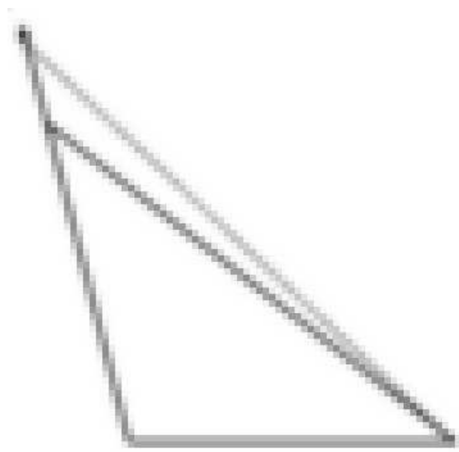

Figure S-b:

Most of these points were made, though less clearly in (Sloman, 1971) which also emphasised the fact that for mathematical reasoning the use of external diagrams is sometimes essential because the complexities of some reasoning are too great for a mental diagram. (These points were generalised in Sloman 1978 Chapter 6). Every mathematician who reasons with the help of a blackboard or sheet of paper knows this, and understands the difference between using something in the environment to **reason** with and using physical apparatus to do **empirical** research, though it took some time for many philosophers of mind to notice that minds are extended. (The point was also made in relation to reference to the past in P.F. Strawson's 1959 book, Individuals, An essay in descriptive metaphysics.)

NOTE ADDED 12 Sep 2012: DIAGRAMS CAN BE SLOPPY

In many cases a mathematician constructing a proof will draw a diagram without bothering to ensure that the lines are perfectly straight, or perfectly circular, etc., or that they are infinitely thin (difficult with line drawing devices available to us). That's because what is being studied is not the particular physical line or lines drawn on paper or sand, etc. The lines drawn are merely representations of perfect Euclidean lines whose properties are actually very different, and very difficult to represent accurately on a blackboard or on paper. E.g. drawing an infinitely thin line has been a problem.

In fact, the lines don't need to be drawn physically at all: they can be imagined and reasoned about, though in some cases a physical drawing can help with both memory and reasoning.

BACK TO CONTENTS

Perception of affordances

All this seems to be closely related to the ability of animals to perceive affordances of various kinds, discussed in Gibson(1979).
I suggest that James Gibson's theory of perception of affordances, is very closely related to mathematical perception of structure, possibilities for change, and constraints on changes (structural invariants). This goes far beyond what he actually wrote about affordances, as far as I know. Gibson's ideas are summarised, criticised and extended
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#gibson

In particular, the kind of mathematical reasoning about infinite ranges of possibilities and implications of constraints, seems to be closely related to the ability of young children and other animals to discover possibilities for change in their environments, and abilities to reason about invariants in subsets of possibilities that can be relied on when planning actions in the environment.

This leads to the notion of a "toddler theorem" discussed in these pages:
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html

I suspect that the reasoning using schematic diagrams illustrated above, and also illustrated in Pardoe's proof of the Triangle Sum Theorem, shares features with animal reasoning about affordances, in which conclusions are reliably drawn about invariants that are preserved in a process, or about impossibilities in some cases -- e.g. it is impossible to completely enclose a bounded area using only two straight lines.
Why is it impossible?

There have been attempts to simulate mathematical reasoning using diagrams by giving machines the ability to construct and run simulations of physical processes. But that misses the point: a computer running a simulation in order to derive a conclusion can handle only the specific values (angles, lengths, speeds, for example) that occur in the initial and predicted end states when the simulation runs. Moreover, the simulation mechanisms have to be carefully crafted to be accurate. In contrast, as pointed out above, a human reasoning about a geometrical theorem does not require precision in the diagrams and the conclusion drawn is typically not restricted to the particular lengths, angles, areas, etc. but can be understood to apply to infinitely many different configurations satisfying the initial conditions of the proof. (See the recent discussion between Mary Leng and Mateja Jamnik, in The Reasoner.)

This seems to require something very different from the ability to run a simulation: it requires the ability to manipulate an **abstract representation** and to interpret the results of the manipulation in the light of the representational function of the representations manipulated. In other words mathematical thinking using diagrams and imagined transformations of geometrical structures, as illustrated above, inherently requires meta-cognitive abilities to notice and reason about features of a process in which semantically interpreted structures are manipulated. The noticing and reasoning need not itself be noticed or reasoned about, although that may develop later (as seems to happen, in different degrees, in humans).

I suspect that many animals, and also pre-verbal human children have simplified versions of that ability, but do not know that they have it. They cannot inspect their reasoning, evaluate it, communicate it to others, wonder whether they have covered all cases, etc. There seems to be a kind of meta-cognitive development that occurs in humans, perhaps partly as a result of learning to communicate and to think using an external language. It may be that some highly intelligent non-human reasoners have something closely related. But we shall need more detailed specifications of the reasoning processes and the mechanisms required, before we can check that conjecture.
(Annette Karmiloff-Smith's ideas about "Representational Redescription", in "Beyond Modularity" (1992) are also relevant.)

We also need more detailed specifications in order to build robots with these "pre-historic" mathematical reasoning capabilities -- which, as far as I can tell, no AI systems have at present. Unlike Roger Penrose, who has noticed similar features of mathematical reasoning, I don't think there is any obvious reason why computer based systems cannot have similar capabilities. However it may turn out that there is something about animal abilities to perceive, or imagine, processes of continuous change at the same time as noticing logically expressible constraints or invariants of those processes that requires information processing mechanisms that have so far not been understood. Alternatively, it may simply be that no high calibre AI programmers have attempted to implement competences of the sorts required to invent and understand Pardoe's proof, or many of the traditional proofs used in Euclidean geometry.

These ideas suggest a host of possible investigations of ways in which human capabilities change, along with the reasoning competences of intelligent animals such as squirrels, elephants, apes, cetaceans, octopuses, and others.

BACK TO CONTENTS

The perpendicular stretch theorem (The need for case analysis in some proofs)

The Median Stretch Theorem (MST above), and the Side Stretch Theorem (SST) on which it depends, both require a single diagram. Distortions of the diagram may produce new figures that look different but they do not require any new form of reasoning.

However there are some theorems in Euclidean geometry whose proof requires use of more than one diagram, because the theorem has a kind of generality that covers structurally different cases. An example of such a theorem is a proof that the area of a triangle is half the area of a rectangle with the same base length and the same height: Area = 0.5 x Base x Height. The reason for requiring more than one diagram (unless there is a proof I have not encountered) will be explained below.

A non-diagrammatic algebraic proof may be possible using the Cartesian-coordinate based representation of geometry, but that is not what this discussion is about.

It is highly regrettable that our educational system produces many people who have simply memorised the Area formula, without ever discovering a proof or being shown one, or even being told that there is a proof, though some may have done experiments weighing triangular and rectangular cards. I shall try to explain how this formula could be proved, though I'll expand the usual proof to help bring out differences between this theorem and previous theorems, explaining why this theorem requires different cases to be dealt with differently.

Consider a theorem related to Figure P-a below, which is subtly different from Figure M (a), above.

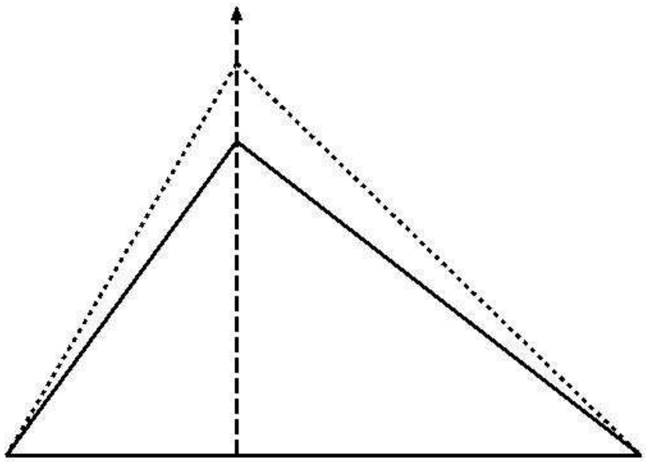

Figure P-a:

Figure P-a includes a straight line drawn between a vertex of the triangle and the opposite side, extended beyond the vertex as indicated by the dashed arrow. In figure M the line used was a **median**, joining the mid-point of a side to the opposite vertex. Here the line is not a median but is **perpendicular** to the opposite side. (In some cases the median and the perpendicular are the same line. Which cases?)

You should find it **obvious** that if the top vertex of the triangle with solid black sides shown in Figure P-a, above, is moved further away from the opposite side (the base), along a line perpendicular to the opposite side (the dashed arrow), then the area enclosed by the triangle must increase. This could be called the "Perpendicular Stretch Theorem" (PST), in contrast with the "Median Stretch Theorem" (MST), which used a line drawn from the middle of the base.

In this figure it is obvious that moving the vertex up the perpendicular will produce a new triangle that encloses the original one. Figure P-a shows why it is obvious, though the Side Stretch Theorem shown in figure S, above, could used to prove this, by dividing the figure into two parts, just as it was used to prove the Median stretch theorem. (As with MST, there is a corresponding theorem about the area decreasing if the vertex moves in the opposite direction on the perpendicular.)

But there is a problem, which you may have noticed, a problem that did not arise for the median stretch theorem. The problem is that whereas any median from the midpoint of one side to the opposite vertex will go through the interior of the triangle, the perpendicular from a side to the opposite vertex may not go through the interior of the triangle, a problem portended by part (b) of Figure M.

BACK TO CONTENTS

A problem with the proof using Figure P-a

Observant readers may have noticed that the reasoning based on Figure P-a has a flaw, since not all movements of a vertex of a triangle perpendicularly away from the opposite side will produce a new triangle that encloses the original one: for example if one of the interior angles (e.g. the one in the left in Figure P-b, below) is obtuse (greater than a right angle), so that the top vertex does not start off perpendicularly above the base of the triangle. The line perpendicular to the "base" that passes through the vertex need not pass through the base, though it will pass through a larger line extending the base, as shown in Figure P-b, which is derived from Figure P-a, by shifting the upper vertex over to the left, so that the perpendicular indicated by the dotted arrow moves outside the triangle, and no longer intersects the base (the side opposite the vertex under consideration), though it intersects the line extending the base.

Figure P-b:

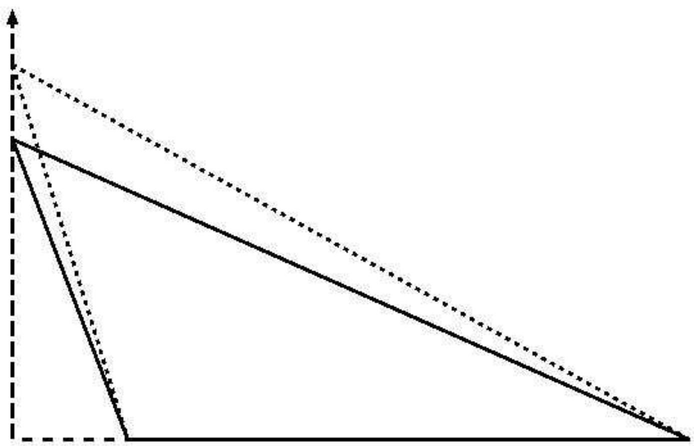

In this case, moving the top vertex upwards will not produce a new triangle enclosing in the old one, because one of the sides of the triangle will move so as to cross the triangle, as illustrated in Figure P-b. So now the proof that the area increases cannot be based on containment: the new triangle produced by moving the vertex upward does not include the old triangle, as in the previous configuration. Is there a way of reasoning about this new configuration so as to demonstrate an invariant relation between direction of motion of the vertex and whether the area of the triangle increases or decreases?

Some readers may notice a way of modifying the proof to deal with figure P-b, thereby extending the proof that moving the vertex further from the line in which the opposite side lies, always increases the area. It is an extension insofar as it covers more cases. Of course, the original proof covered an infinite set of cases, but that infinite set can be extended.

A clue as to how to proceed can come from considering how to prove that moving the vertex of a triangle parallel to the opposite side, as illustrated in Figure P-c, below, cannot change the area.

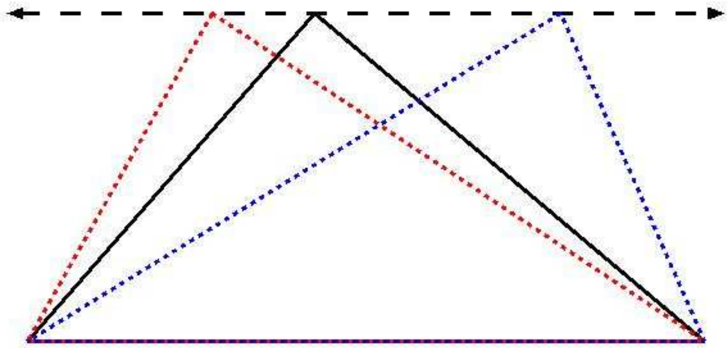

Figure P-c:

TO BE CONTINUED

I shall later extend this discussion by showing how to relate the area of a triangle to the area of a rectangle enclosing it. It will turn out that the triangle must always have half the area of the rectangle, if the rectangle has one side equal in length to a side of the triangle and the other side equal in length to the perpendicular of the triangle. Proving this requires dealing with figures P-a and P-b separately.

The proof using a rectangle requires introducing a new discontinuity into the configuration: dividing up regions of the plane so that they can be compared, added, and subtracted.

Some readers will be tempted to prove the result by using a standard formula for the area of a triangle. In that case they first need to prove that the formula covers all cases, including the sort of triangle shown in Figure P-b.

For anyone interested, here's a hint. Consider Figure TriRect, below. Try to prove that every triangle can be given an enclosing rectangle, such that every vertex of the triangle is on a side of the rectangle and two of the vertices are on one side of the rectangle, and at least two of the vertices of the triangle lie on vertices of the rectangle.

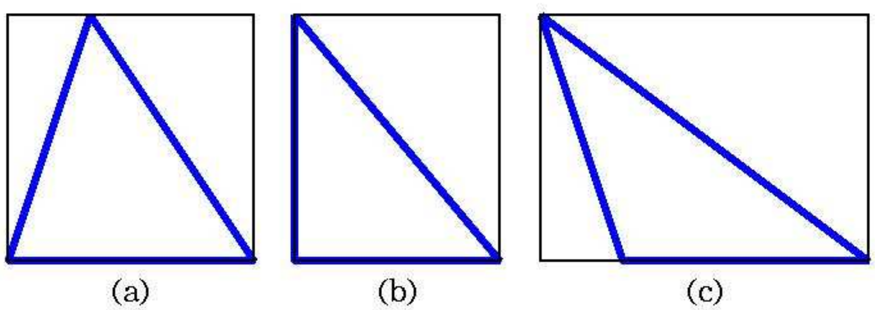

Figure TriRect:

Can you prove something about the area of a triangle by considering such enclosing rectangles?

Many mathematical proofs are concerned with cases that differ in ways that require different proofs, though sometimes there is a way of re-formulating the proof so that the same reasoning applies to all the structurally distinct cases. A fascinating series of examples from the history of mathematics is presented in (Lakatos, 1976)

NOTE
A hard problem for human and animal psychology, and studies of evolution of cognition, is to explain how humans (and presumably some other animals capable of intelligent reasoning about their affordances), are able to perform these feats. How do their brains, or their minds (the virtual information-processing machines running on their brains), become aware that the special case being perceived shares structure and consequences of that structure, with infinitely many other configurations, the majority of which have never before been seen or thought about.

For an explanation of the notion of a virtual machine made of other concurrently active virtual machines, some of which also interact with the environment, see
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html

BACK TO CONTENTS

Toward robot mathematicians discovering geometry

It will be some time before we have robot mathematicians that understand Pardoe's proof, or the proofs of the 'Stretch' theorems summarised above (Median stretch, Side stretch, Perpendicular stretch theorems), or can think about how to compute the area of a triangle, or can discover the existence of prime numbers by playing with blocks (in the manner described here), or can perceive and make use of the many different sorts of affordance that humans and other animals can cope with (including, in the case of humans: proto-affordances, action affordances, vicarious affordances, epistemic affordances, deliberative affordances, communicative affordances), many described in this presentation on Gibson's theories.

Even longer before a robot mathematician spontaneously re-invents Pardoe's proof? (Or the proofs in Nelsen's book.)

For some speculations about evolution of mathematical competences see

- http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#mathcog If learning maths requires a teacher, where did the first teachers come from?
- http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler Why (and how) did biological evolution produce mathematicians?
- http://www.cs.bham.ac.uk/research/projects/cosy/papers#tr0802 Kantian Philosophy of Mathematics and Young Robots (MKM08)
- http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0807 The Well-Designed Young Mathematician (AI Journal 2008)

Chemical computation A deeper question is whether there is something about the information-processing engines developed and used by evolution that are not modelled in turing machines or modern computing systems, or have totally intractable complexity on Turing machines or modern computers. I shall later produce some speculative notes on whether there are deep differences between chemistry-based computation and more familiar forms of computation.

If there are differences I suspect they may depend on some of the following:

- Chemical processes involve both continuous changes in spatial and structural relations and also the ability to cross a phase boundary and snap into (or out of) a discrete stable state that resists change by thermal buffeting and other processes. This stability could rely on quantum mechanisms.
- They also allow multiple constraints to be exercised by complex wholes on parts, which allow certain forms of motion or rotation or chemical behaviours but not others.
- Some switches between discrete states, or between fixed and continuously variable states can be controlled at low cost in energy by catalytic mechanisms.

It is clear that organisms used chemical computation (= chemical information processing) long before neural mechanisms were available. Even in organisms with brains, chemical information processing persists and plays a fundamental role (e.g. building brains and supporting their functionality). It is possible that more information will emerge later from the Meta-Morphogenesis project Sloman(2013).

BACK TO CONTENTS

Comparison with logical proofs

Many mathematical proofs involve sequences of logical formulae or equations, with something altered between stages in the sequence. Those sequences can be thought of as processes, but they are essentially discrete, discontinuous processes. For example, consider the transformation from (P1) and (P2) to (C) in this logical proof:

Premisses
(P1) All Humans are Mortal (or (All x)H(x) -> M(x))
AND
(P2) All Greeks are Humans (or (All x)G(x) -> H(x))

Conclusion
(C) All Greeks are Mortal (or (All x)G(x) -> M(x))

For someone who does not find this obvious, the proof can first be transformed into a diagram which initially represents (P1), then adds the information in (P2), then shows how that includes the information in (C), showing the proof to be valid.

This can be thought of as a process, but the steps are distinct and there are not meaningful intermediate stages, e.g in which the antecedent "H(x)" and the implication arrow "->" are gradually removed from the original implication, and the word "Socrates" gradually replaces the variable "x". Nevertheless the proof can be expressed diagrammatically using Euler Circles as in Figure Syll (often confused with Venn Diagrams, which could also be used).

Figure Syll: 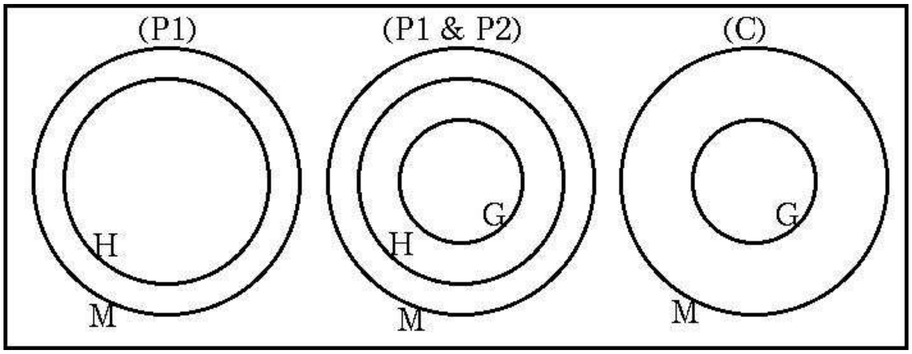

In (Sloman, 1971) I suggested that both types of proof could be regarded as involving operations on representations that are guaranteed to "preserve denotation". This is an oversimplification, but perhaps an extension of that idea can be made to work.

In the Pardoe proof, "preserving denotation" would have to imply that a process starting with the initial configuration in Figure Ang3, and keeping the triangle unchanged throughout, could go through the stages in the successive configurations depicted, without anything in the state of affairs being depicted changing to accommodate the depiction, apart from the changes in position and orientation of the arrow, as shown.

This implies, for example, that there are no damaging operations on the material of which the structures are composed. (I suspect there is a better way to express all this.)

Cathy Legg has presented some of the ideas of C.S. Peirce on diagrammatic reasoning in (Legg 2011) It is not clear to me whether Peirce’s ideas can be usefully applied to the kinds of reasoning discussed here, which are concerned with geometrical reasoning as a biological phenomenon with roots in pre-human cognition, and properties that I suspect could be replicated in robots, but have not yet been, in part because the phenomena have not yet been understood.

BACK TO CONTENTS

A partial list of references (to be expanded)

- http://math.berkeley.edu/~rbayer/09su-55/handouts/ProofByPicture-printable.pdf
  Robertson Bayer, Proof By Picture (PDF lecture slides),
  University of California, Berkeley Math 55, Summer 2009
  (A collection of diagrammatic proofs of mathematical theorems, most of them
  non-geometric -- e.g. geometric proofs of theorems in number theory.
  Includes the ’Chinese’ proof of Pythagoras’ Theorem.)

- **Added 6 Aug 2013**
  http://www.mmrc.iss.ac.cn/~xgao/paper/book-area.pdf
  Shang-Ching Chou, Xiao-Shan Gao and Jing-Zhong Zhang,
  Machine Proofs In Geometry: Automated Production of Readable Proofs for Geometry Theorems,
  World Scientific, Singapore, 1994,

In the Foreword, Robert Boyer writes:

> The stunning beauty of these proofs is enough to rivet the reader's attention into learning the method by heart.
> The key to the method presented here is a collection of powerful, high level theorems, such as the Co-side and Co-angle Theorems. This method can be contrasted with the earlier Wu method, which also proved astonishingly difficult theorems in geometry, but with low-level, mind-numbing polynomial manipulations involving far too many terms to be carried out by the human hand. Instead, using high level theorems, the Chou-Gao-Zhang method employs such extremely simple strategies as the systematic elimination of points in the order introduced to produce proofs of stunning brevity and beauty.

It seems that hundreds of geometrical theorems have been proved, many of them non-trivial, and the theorem prover translates the proofs into latex suitable for generating human-readable versions!
One of the key ideas is a representation of triangles (and other polygons) that allows complex case analysis on traditional geometric reasoning to be collapsed into a single case, or a very small number of cases. However the process by which such a form of representation is discovered and understood to be adequate for the purpose is not modelled by the theorem prover. (That's not a criticism of what seems to be an outstanding achievement.) Moreover, the mechanisms and form of representation assume, as far as I can tell, a more sophisticated grasp of geometry than a young child or a non-human animal seems to need for discovering and reasoning about invariant properties of spatial processes and the affordances involved. Nevertheless, it may turn out that the forms of representation and reasoning in this book (which I have not yet understood in detail) may provide useful clues regarding more primitive, yet powerful, abstractions available to our ancestors, pre-verbal children, and other species.

- Kenneth Craik, **The Nature of Explanation,**
  Cambridge University Press, 1943, London, New York,

  NOTE:
  Craik proposed that biological evolution produced animals with the ability to work out what the consequences of an action would be without performing the action, by making use of an abstract **model** of the situation in which the action is performed. It is not clear to me that he noticed the difference between
  - running a detailed model of a specific situation to discover the specific consequences, which some current AI systems (e.g. game-engines) can do, and
  - noticing an invariant property of such a process with different starting configurations as required for understanding why a strategy will work in a (possibly infinite) class of cases.

  I think he came close, but did not quite get there, but I have read only the 1943 book.

- Susan Blackmore, **Zen and the Art of Consciousness,** Oneworld Publications, 2011
  (Previous title: Ten Zen Questions,)
  Discussed in:
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/blackmore-zen-consciousness.html

- George B. Dyson,
  Darwin Among The Machines: The Evolution Of Global Intelligence,
  Addison-Wesley, 1997,

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

- **Entertaining video tutorial on Pythagoras and irrationality**
  showing the connection between irrationality of square root of 2, and Pythagoras' theorem.
  https://www.youtube.com/watch?v=X1E7I7_r3Cw
  "What's up with Pythagoras", by Vi Hart. See more of her work:
  https://www.youtube.com/user/Vihart
  And here: http://vihart.com
  You may have to pause and replay bits of her videos, to take in some of the details.

- Immanuel Kant, **Critique of Pure Reason,** 1781,
  Translated (1929) by Norman Kemp Smith, London, Macmillan,

- Annette Karmiloff-Smith, 1992,
  Beyond Modularity: A Developmental Perspective on Cognitive Science,
  MIT Press,
  Partly summarised, and discussed in:
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html

- http://adrenaline.ucsd.edu/Kirsh/Articles/Earwig/earwig-cleaned.html
  David Kirsh, 1991
  "Today the earwig, tomorrow man?,
  in **Artificial Intelligence,** 47, 1, pp. 161--184,

- Imre Lakatos, 1976, **Proofs and Refutations,** CUP, Cambridge, UK,

- Roger B. Nelsen, (Often mis-spellt as "Nelson")
  Proofs without words: Exercises in Visual Thinking,
  Mathematical Association of America, Washington DC, 1993,

- http://hdl.handle.net/10289/5153
  Catherine Legg, What is a logical diagram?
  Paper presented at Mini-Conference on Logical Pragmatism,
  Auckland, New Zealand, February 25, 2011.

- **Mary Leng interviews Mateja Jamnik on Spatial Reasoning**,
  In **The Reasoner** 7(1), Jan 2012. pp 1--4.
  http://www.kent.ac.uk/secl/philosophy/jw/TheReasoner/vol7/TheReasoner-7(1).pdf

- J. Sauvy and S. Sauvy, **The Child's Discovery of Space: From hopscotch to mazes -- an introduction to intuitive topology,**
  Penguin Education, 1974, Translated from the French by Pam Wells,

- http://www.cs.bham.ac.uk/research/projects/cogaff/04.html#200407
  Aaron Sloman, 'Interactions between philosophy and AI:
  The role of intuition and non-logical reasoning in intelligence',
  Proc 2nd IJCAI, 1971, London, pp. 209--226, William Kaufmann,
  Also reprinted in AI journal, 1971, and in "The Computer Revolution in Philosophy", Chapter 7, 1978)
  http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap7
  That article was originally written as a critique of McCarthy and Hayes 1969:
  http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html
  "Some philosophical problems from the standpoint of Artificial Intelligence"
  (In **Machine Intelligence 4**)

- http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#15
  Aaron Sloman, Actual Possibilities, in
  Principles of Knowledge Representation and Reasoning: Proc. 5th Int. Conf. (KR '96),
  Eds. L.C. Aiello and S.C. Shapiro, Morgan Kaufmann, Boston, MA, 1996, pp. 627--638,

- http://www.cs.bham.ac.uk/research/projects/cogaff/misc/fully-deliberative.html
  Aaron Sloman, 2006,
  Requirements for a Fully Deliberative Architecture (Or component of an architecture),
  Research Note: COSY-DP-0604, School of Computer Science, University of Birmingham, UK,

- Aaron Sloman, 2013, Virtual Machinery and Evolution of Mind (Part 3):
  Meta-Morphogenesis: Evolution of Information-Processing Machinery, in *Alan Turing - His Work and Impact,* Eds. S. B. Cooper and J. van Leeuwen, Elsevier, Amsterdam, pp. 849-856,
  http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d
  Ongoing work on the Meta-Morphogenesis project is available here:
  http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html
  and an early slide presentation on Meta-Morphogenesis is here:
  http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk107

- http://dspace.mit.edu/handle/1721.1/6894
  Gerald. J. Sussman,
  A computational model of skill acquisition, American Elsevier, 1975,

- Max Wertheimer, **Productive Thinking**
  New York, Harper and Brothers, 1945

Acknowledgements

See the acknowledgements section of the paper on P-Geometry
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/p-geometry.html#acknowledge

Offers of help in making progress will be accepted gratefully, especially suggestions regarding mechanisms that could enable robots to have an intuitive understanding of space and time that would enable some of them to re-discover Euclidean geometry, including Mary Pardoe's proof.

I believe that could turn out to be a deep vindication of Immanuel Kant's philosophy of mathematics. Some initial thoughts are in my online talks, including:

http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler
Why (and how) did biological evolution produce mathematicians?

Related documents

A discussion of the Triangle Sum Theorem, especially Mary Pardoe's proof, previously part of this file has been moved to:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html

A partial index of discussion notes in this directory is in
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html

See also this discussion of "Toddler Theorems":
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html

Discussion of various proofs of the Triangle Sum Theorem (Angles):
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html

Discussion of possibility of adding rigid motion to Euclidean geometry
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/p-geometry.html

This document illustrates some points made in a draft, incomplete, discussion of transitions in information-processing, in biological evolution, development, learning, etc. here.

That document and this one are both parts of the Meta-Morphogenesis project, Sloman(2013), partly inspired by Turing's 1952 paper on morphogenesis.

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

More general overviews of relations between biology, evolution, mathematics and philosophy of mathematics are presented in these incomplete drafts:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/mathsem.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html

---

Document history

Installed: 9 Sep 2012
28 May 2013 moved section on Triangle Sum Theorem to separate file.
Please report bugs (A.Sloman@cs.bham.ac.uk)

Last updated:
6 Aug 2017: Made connection with perception of impossibilities and affordances clearer.
9 Apr 2017: Added disclaimer below.
8 Nov 2016: Further reformatting -- and a few notes added.
14 Aug 2015 Reformatted, with some minor changes of content.
20 Dec 2013 minor changes.

2 Sep 2013: rearranged figure M
20 Aug 2013: fixed broken references to Pardoe's proof of the triangle sum theorem.
6 Aug 2013 (added ref to Chou et al. 1994)
27 May 2013 (more on non-metrical theorems/proofs); 29 May 2013;
(Reorganised March 2013, 7th May, 29th May)
24 Sep 2012; Oct ; Nov ; Dec ; Jan 2013; Feb 2013;

Installed and maintained by Aaron Sloman
School of computer science, Univ. of Birmingham: http://www.cs.bham.ac.uk/

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