

If learning maths requires a teacher, where did the first teachers come from?

or

Why (and how) did biological evolution produce mathematicians?

Aaron Sloman, School of Computer Science, UoB.

<http://www.cs.bham.ac.uk/~axs>

Slides here

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#mathcog10>

This is the latest in a series of presentations on this topic, under various titles.

This title was used at the IR-Lab meeting, Computer Science, Univ of Birmingham 26th Jan 2010

and a talk at Learning Sciences Research Institute, University of Nottingham 2nd Feb 2010.

See also the paper with this title for AISB2010 symposium.

<http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1001>

(X) Apologies and Notes

Apologies/warnings:

- My work straddles so many disciplines that I cannot keep up with most of what has been written that is relevant.
I welcome pointers to things I should have read or known about.
- My slides are too cluttered for presentations: I write them so that they can be read by people who did not attend the presentation.
So please ignore what's on the screen unless I draw attention to something.

SHORT (BUT COMPLEX) ANSWER:

Human mathematical competences are a by-product of evolutionary pressures

- to cope with a complex, changing, 3-D environment
- with rich interacting structures and processes,
- in which animals need to produce, prevent, or maintain various states and processes,
- which cannot be dealt with adequately using only learning based on sensorimotor statistics: an **exosomatic** ontology is needed
- and mechanisms for manipulating representations of possible processes (proto-affordances)

SHORT (BUT COMPLEX) ANSWER:

Human mathematical competences are a by-product of evolutionary pressures

- to cope with a complex, changing, 3-D environment
- with rich interacting structures and processes,
- in which animals need to produce, prevent, or maintain various states and processes,
- which cannot be dealt with adequately using only learning based on sensorimotor statistics: an **exosomatic** ontology is needed
- and mechanisms for manipulating representations of possible processes (proto-affordances)

We can try to understand this by trying to design robots that cope in such a world.

But first we need a much better understanding of **the requirements**.

What are the problems that animals of various sorts (including young humans) have to be able to deal with? (including acting, thinking, thinking about thinking...)

Kant on Synthetic A priori Knowledge (e.g. mathematical knowledge).

Immanuel Kant wrote:

“There can be no doubt that all our knowledge begins with experience.

For how should our faculty of knowledge be awakened into action did not objects affecting our senses partly of themselves produce representations, partly arouse the activity of our understanding to compare these representations, and, by combining or separating them, work up the raw material of the sensible impressions into that knowledge of objects which is entitled experience?

In the order of time, therefore, we have no knowledge antecedent to experience, and with experience all our knowledge begins.”

Critique of Pure Reason

How can experiences of how the world is be a trigger for acquiring knowledge whose truth does not depend on how the world is?

Some examples will be presented and then a schematic, incomplete, theory.

Help completing it will be very welcome.

The very short answer

- Some species, that are able to learn empirical generalisations, including humans,
- are genetically predisposed to seek a more economical **generative** representation of what has been learnt (a “framework theory”?)
- that both makes what was previously learnt **derivable** and therefore explainable,
- and permits well-founded reasoning about complex novel situations.

It is well known that something like this happens during language learning, e.g. the switch from learning patterns to using a generative syntax.


I conjecture that that evolved later, as a special case of something more general found in human and some non-human species.

Syntax of natural languages is special because it includes so many exceptions to the main rules, requiring a complex extension to the architecture.

Hence the phenomenon of “U-shaped” language learning.

Embodiment theories: fashionable but confused


Often assume close coupling between mind/brain and environment through body, sensors, and effectors.



A multi-stable dynamical system closely coupled with the environment through sensors and effectors.

All the semantics may be **somatic** referring only to states of sensors and effectors and their relations.

Evolution produced something more sophisticated



A multi-layered, multi-component dynamical system with most sub-systems not coupled to the environment, but some able to refer to things in the environment that cannot be sensed, e.g. past, future, remote, and hypothetical entities. This uses an **exo-somatic** ontology.

But there are many questions about how the external environment is represented.

E.g. a unified global, fully metric coordinate system?

Compare large numbers of topological and semi-metrical information structures, some of them dynamically changing.

Main points (really needs several hours)

- Much mathematical knowledge is of interest because it is applicable to the world
- Much of it is discovered through first discovering general truths empirically then re-construing them as non-empirical
- Even young children can do this, but they don't know they are doing it. (Examples later)
- This often involves developing new forms of representation, and new ontologies, used in new theories, from which previous discoveries turn out to be derivable – a new status: **necessary** or **semi-necessary**, depending on the type of theory.
- Some of these are **framework theories**: we can't do without them. (A problematic concept.)
- At first the new knowledge is merely usable (not explicitly addressable).
- At a later stage **explicit**, but **fragmentary**, articulation of the knowledge develops.
- At a still later stage the articulation becomes **formal** and **systematic**.
- All this depends on both (a) the world having a structure that makes these developments possible and useful, (b) the learner having a layered information processing architecture of a certain sort (growing self monitoring virtual machines)
- Some developments depend on cultural evolution enhancing individual development.
- The process is not infallible (Lakatos) but bugs can be detected and sometimes eliminated (if the architecture supports that). (Examples later.)
- This supports Kant on causation: against Hume (causation=correlation). (Both needed.)

A new way to do philosophy of mathematics

Most philosophers of mathematics try to understand the nature of mathematical knowledge “from the inside” – starting from features of mathematical knowledge and concepts that all mathematicians (hopefully) would recognise and agree on.

An alternative is to see human mathematics as a product of biological evolution of many species, extended and enriched by development of many individuals using the cognitive mechanisms produced by evolution and the environment.

All of that depends on the environment presenting challenges and opportunities, both to evolution and to individuals.

I believe that developing this approach can provide new support for an updated version of Kant’s philosophy of mathematics.

It was the belief about 60 years ago that Kant was basically right about mathematics as I had experienced it (doing a degree in mathematics and physics), whereas most of the philosophers I met (mainly in Oxford) thought Kant’s view had been refuted by Einstein and Hume’s view was the main alternative: namely that mathematical knowledge was non-empirical, as Kant claimed, but essentially “analytic”, or merely a collection of definitional truths and their consequences.

What came first: mathematics teachers or pupils?

It is normally assumed that mathematical knowledge has to be acquired with the help of adults who already know mathematics.

If that were true, mathematical knowledge could never have developed on earth, for there were no mathematics teachers in our early evolution.

These slides outline a multi-disciplinary research programme (using a spiral of theory, experiments, challenges and evidence), on a conjecture about how mathematical competences originally evolved under pressure from biological needs.

Such pressures led to mechanisms, forms of representation and architectures that supported two very different forms of knowledge acquisition

- **Empirical learning**, including use of statistical information, correlations, etc.
- **Working things out**, including designing novel structures and mechanisms, and also making predictions, forming explanations, using not statistical but structural information – e.g. geometry, topology, properties of matter... etc.

Later on, mathematical knowledge became formalised, ritualised, and transformed into an object of study for its own sake, and culturally transmitted: through mathematics teaching.

But that teaching works only because the mechanisms produced by evolution are there in every normal child, though there may be differences leading to different speeds of learning and different preferred directions of learning.

Understanding mechanisms helps us understand products

If we can find a way to understand the mechanisms involved in production of mathematical knowledge, and how they relate to the environment in which mathematical knowledge is developed, this will help us understand some unobvious features of the products of those mechanisms – answering philosophical questions in a new way.

- One way to understand the mechanisms is to look closely at some of the processes they produce (e.g. in young children and some other animals)
- and to try to understand the cognitive functions involved.
- The ideas can in principle be tested by trying to use the proposed mechanisms in designing a working robot using those mechanisms, to check that they can produce the relevant kinds of learning.
- This task is doubly difficult because it is difficult (a) to find out what the mechanisms do and (b) to design and implement machines that do that.
- In particular, the mechanisms seem to be closely related to the functions of visual perception in humans, and we are very far from understanding that: AI vision systems are pathetically inadequate compared with humans.
- This approach turns out to be **very** different from recent work emphasising the role of embodiment in animal and robot cognition, which I think is mostly very confused because it presents metaphors as mechanisms.

Sample of views on the nature of mathematical knowledge

- David Hume:
All knowledge is either empirical or relations between ideas (analytic) or “sophistry and illusion”.
- J.S. Mill: Mathematical knowledge is empirical (that’s how much of it starts)
- Immanuel Kant:
Mathematical knowledge is synthetic (non-analytic), non-empirical and necessary.
But he allows that discoveries are triggered by experience.
- Logicism of Russell and Frege:
All mathematical knowledge is just logical knowledge
(Frege - arithmetic Yes, but geometry No)
- Bertrand Russell (1918): [Mathematics may be defined as the subject in which we never know what we are talking about, nor whether what we are saying is true.](#)
- Hilbert: Mathematical knowledge (of the infinite?) is about formal systems.
- Brouwer, Heyting, & others: intuitionists, finitists.
- Wittgenstein (RFM) (a form of conventionalism? Anthropologism?)
- Lakatos: Mathematics is Quasi-empirical (only half right?).
- Richard Feynman: [the language nature speaks in.](#) (Yes. Sort of – see below).

These are shallow and inadequate summaries, to illustrate the variety of philosophies of mathematics.

A new proposal: Robotics-based philosophy of mathematics

Kant (*Critique of Pure Reason*(1781)) claims that mathematical knowledge starts from experience of **how things happen to be** and then somehow becomes non-empirical knowledge about **what has to be**.

Even he did not put it as baldly as that.

This presentation gives examples of things learnt empirically by infants and toddlers (and future robots) that later change their epistemological status, in some of them.

However, it is not yet clear what exactly changes, nor what forms of representation, mechanisms and architectures are involved in making this happen.

Providing a theory that shows how future human-like robots could go through such transitions will provide a new vindication for Immanuel Kant's view of mathematical knowledge as (a) synthetic, (b) necessarily true, (c) non-empirical (but NOT innate).

But things are more complex and varied than he realised, as I'll try to show.

To gain deep understanding, don't study just one species

Trying to do AI or cognitive science by developing only **one** type of system, or studying just one species

(e.g. “human level AI”)

is like trying to do physics by investigating only how things behave near the leaning tower of Pisa.

Or studying only the motion of planets.

we need to understand spaces of possibilities

and the tradeoffs between alternative designs: so **look at different species**.

Don't assume all effective information-processing mechanisms have to be rational

(as in Dennett's “intentional stance”, Newell's “knowledge level”.)

Evolution provided useful reflexes: reflexes are neither rational nor irrational.

Some useful reflexes are cognitive

That can include reflexes that generate goals/motives.

Not all goals are chosen because the individual knows associated rewards/costs/benefits:

See

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html>

Architecture-based motivation vs reward-based motivation.

A collaborative multi-disciplinary research project

Fully developing and testing the ideas presented here
and working out all their implications
requires a team of researchers

- investigating varieties of learning and problem solving in humans and other animals
- including patterns in the way they develop
- asking questions that are not normally asked,
e.g. about the developing **ability to work things out**
- trying to understand how these patterns evolved
- and what mechanisms have to be genetically determined
- and how the genome determines them
- and how all this might be implemented in robots, e.g.
 - which forms of representation
 - which information processing mechanisms
 - which architectures (including self-monitoring, self-extending architectures)
 - which types of development of the architectureare capable of producing the required phenomena.
- and what it tells us about requirements for being a mathematician
- and the implications for various questions in philosophy of mathematics

What is empirical information?

Empirical information (not just beliefs), e.g.

- about specific situations,
- about general truths,
- about associations,
- about how to do something,
- about what will happen in certain circumstances,

can be **acquired** in various ways, e.g. by:

- observing something that happens, possibly as a result of performing some observation, introspection, measurement or manipulation (in some cases repeatedly),
- being informed by another individual, newspapers, a book, a dream, ...
- derivation from other information, using logic, mathematics, statistics, analogy, etc.
- guess-work, hunches
- from the genome (**Genetic inheritance is the main source of general information in most species.**)

What makes many kinds of information empirical (as Popper noted) is not **the source**, but **the possibility of observations and experiments contradicting the information**:

If a perceived counter-example could turn up, a prediction could fail, a situation could be observed to have been previously misperceived, or a method could sometimes not work, etc. (as often happens in science) the information must then be empirical. (**Finding exactly which portions are empirical may be tricky.**)

I. Lakatos in *Proofs and Refutations*: showed that mathematics can be “quasi empirical”.

The counter-examples in mathematics can occur in thought, not just perceived phenomena.

Theories about what is **possible** (e.g. formation of neutrinos) can only be **confirmed** empirically, not **refuted**. (This needs more explanation: See Chapter 2 of *The Computer Revolution in Philosophy*)

The ubiquity of “toddler theorems”

- I shall give several examples of kinds of empirical discovery that can be transformed into non-empirical discoveries, but my examples merely scratch the surface:
I think there are far, far more toddler theorems of many kinds than anyone has dreamed of (except maybe Piaget?)
- Although there are many examples, I do not believe that anyone in AI, cognitive science, psychology, neuroscience, ... has a theory of how the transformation process works, and what its full scope and limitations are.
- In particular, very many of the discoveries are concerned with spatial structures, processes, and interactions: and AI systems don't yet have the ability to represent and reason about them in the right way (at the right level of abstraction).
Something between simulation and logical/algebraic representation.
- We also don't know how brains can support such processes.
- The processes are not all unique to humans, since creative problem solving observed in some birds and mammals (and octopuses?) needs some of the mechanisms.
That's why I am collaborating with two biologists interested in animal cognition, Jackie Chappell and Susannah Thorpe.
- NB: the recent emphasis on embodiment is seriously confused and misleading.
What's important about being embodied is the nature of the environment and the problems it poses: many different forms of embodiment are compatible with finding solutions. (Sloman 2009)
(But perhaps the virtual machinery required shows convergent evolution: mammal, corvid, octopus.)

Examples

There are many ways in which human children (and perhaps members of some other species) can learn things by exploring the world, which, although learnt empirically have the potential to be reconstrued as non-empirical – either absolutely necessarily true, or at least derivable from a deep and general theory.

- The examples differ considerably in the details of what is learnt, how it is learnt and the information structures involved.
- I shall merely make a few comments on each example, though they all require much more detailed study – including investigation of ways of giving the competences to robots.
- I suspect these examples merely scratch the surface of what goes on in human learning about the environment in childhood.
- After giving examples, I introduce the notion of a [framework theory](#) and try to show that there are two sorts of framework theory.
- The facts discovered empirically alter their status when it is found that they are derivable from a framework theory.
- The precise nature of framework theories and how they differ from one individual to another is left unanswered.

Example: Toddler Steerage

The pictures of babies/toddlers shown in these slides are snapshots from videos available [here](http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/vid/): <http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/vid/>

At work with a broom (about 15 months)



He notices that the broom handle has restricted movement, turns, and uses

Toddler theorem 1:

If an object is between two rigid, rigidly located, vertical bars its sideways motion in the plane of the bars is restricted.

Toddler theorem 2:

If the object is moved perpendicular to the plane of the bars far enough, the restriction is removed.

Later he uses other toddler theorems (a) to deal with the skirting board restricting forward motion of the broom by switching from pushing to pulling and (b) to get the broom to move into the room on the right by rotating the broom so that it points through the doorway.

I am not claiming that at that age a toddler understands (a) that he is using such general principles (b) that he appreciates their mathematical character.

More Examples

Show videos

- Parrot – apparently understanding implications of grasp location when scratching the back of its head with a feather.
- Betty the genius New Caledonian crow.
- Pre-verbal children reasoning about spatial (and social) problems.
- Collecting empirical information with a piano
- Some other examples – see
`http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/vid/`
- Contrast impressive robots that behave but can't (yet) develop a child-like understanding of what they are doing.

E.g. Factory assembly robots, and (as far as I can tell) the amazing BigDog, at Boston dynamics.

Some infant and toddler examples

It takes a child some time to learn how to align a cup so that it will go into or stack on another cup.

Toddler theorem: If cup A is to be inserted into cup B, then the orientations of the cups must be aligned, A must be on the open side of B, and the axis of cup A must point into cup B, but rotation of A about that axis makes no difference.



Inserting picture pieces into recesses.

At a certain age a child can lift a puzzle piece out of a recess, without being able to restore it to the recess, despite being able to put it in the right location.

The problem does not seem to be lack of control of movements.


It is lack of ability to understand the required relationships, possibly because the child's ontology is not yet rich enough to include all the relevant entities and relationships (e.g. boundary, alignment).

So some information available in the optic array is not acquired, represented and used by the child.

Later it is used

What changes in the child's brain, or mind?

Toddler theorem: Inserting an asymmetric picture into its matching recess requires alignment of boundaries of picture and recess.



A Process theorem: Toddler Nesting

Understanding how to stack cylinders.

Consider being faced with a collection of cylindrical cans or mugs of different heights and widths all open at the top, as shown in the upper picture.

What do you have to do to get them all tucked together compactly as in the second picture?


What steps should you go through?

A very young child may try putting one mug into another at random, and then get frustrated because situations result like the one depicted at the bottom, where the top mug cannot be pushed down into the mug below it, as required.

By trial and error a child may learn a sequence of actions that works, e.g.

- (i) put mug 2 into mug 1
- (ii) put mug 4 into mug 3
- (iii) lift mug 3 with mug 4 inside it and put both into mug 2

What if there are five mugs, or nine mugs?



Towards a Toddler Nesting Theorem

Later, through trial and error the child may learn to start at one end of the row, and work across the cylinders systematically, putting each one into a wider one next to it, starting with the narrowest.

Another option involves starting at the other end and working systematically.

But what will happen if the cylinders are not initially arranged in increasing order of width?

Eventually a child may discover that there is a simple stacking rule that will work for any initial configuration, always getting the cylinders compactly stacked.

At that stage, or some time later, many children seem to realise that some of the things they have learnt are **guaranteed** to produce the right result, e.g. no matter how big the mugs are, how they are initially arranged, what they are made of, what colours they are painted, etc., **as long as certain conditions are satisfied**.

Likewise they may indicate that they realise something is **guaranteed** to fail: e.g. they may call you silly for trying: are they just using a very high probability estimate?

This is actually a rule about operations that preserve orderings, and can be relied on without testing, as long as the objects involved keep their size and shape neither disappear nor appear out of nowhere. **Do YOU believe it?** If so, why?

Towards a theory of ordering processes

A child may, after a while, notice patterns in the **processes** that occur during attempts to get containers nested, leading eventually to the following discovery somehow represented in the child, though not necessarily in words or a logical formalism:

- If the container that is widest is moved to one side, then repeatedly the widest remaining container is inserted in it, eventually all the containers will be “tucked together compactly” as in the earlier picture.

A child need not use the word “widest” but may learn to do visual comparisons between containers, and repeatedly discard items after finding something wider.

- Alternatively, starting with the narrowest then repeatedly looking for the narrowest remaining container and inserting the already nested containers into it will achieve the same result (in the reverse order).

This requires a technique for identifying the narrowest remaining container, which can also be done using visual comparisons.

The insertion process will be more complex, as it will require the whole nested collection of previously stacked containers to be lifted and inserted into the next container selected.

- A child may discover, perhaps initially by accident, that a way to avoid repeated searches for the biggest or smallest remaining container, is to start by arranging them in a row in order of size. Then, whether looking for the smallest or the biggest remaining container, the child need only look at one end of the row.

Later, a child may learn that the stacking process can occur in many different orders.

Knowing how, knowing that, and knowing why

The hypothesis being explored here is that in at least some cases, and the container nesting process might be such a case, the learner acquires different sorts of information in the following order

(though older, more sophisticated learners will learn such things in different orders).

- Through trying various things out a child may learn a way of doing things, e.g. by acquiring chains of associations of the form **If the goal is G and the context is C do A**
Compare Nilsson's Teleoreactive programs.
- Later, reflective and communicative competences develop so that the child is able
 - (a) to notice that a certain type of goal can be achieved in a particular way,
 - (b) to enable others to achieve such goals, by observing them and noticing problems, helping them, correcting them, telling them what to do
 - (c) and eventually to formulate one or more strategies that succeed in achieving that type of goal in a variety of contexts.

Still later the child may come to understand **why the strategy works**.

This may at first involve grasping **perceptual invariants** of the process used, such as:

- If items are repeatedly taken from one end of a row ordered by size, what is left is a row ordered by size
- If items are repeatedly taken from a row ordered by **decreasing** size then each can be inserted in the previously taken item
- After each insertion the row will shrink, until no item remains in the row.
- The previous process will end with all the items except the first one, inserted, with nothing left to insert.
- These features of the process do not depend on the colour, material, size, weight of the objects, nor whether it is day time night time, or hot or cold in the room: many features of the context are irrelevant.

Understanding why a procedure must work

Young children have to

1. extend their ontology to include new categories of relationship between the cups (e.g. 'next bigger', 'next smaller')
2. and also so as to develop new categories of process fragment (e.g. 'insert next smaller'),
3. and ways of composing process fragments (e.g. 'repeat ... until ...'), including notions of 'how to start', 'stopping condition', etc.

Doing all that (consciously or not!) can develop


- the competence that enables the task to be completed,
- but not yet the competence to realise that it **must** always work no matter how many cups there are (under certain conditions).

That's what computer scientists do when 'proving correctness'.

- Of the formulations proposed above, the logical one most clearly allows us to compose a procedure for stacking the cylinders, and then **to prove that the procedure achieves the desired goal** on the basis of a collection of assumptions about the objects, e.g. they do not divide or merge or move around of their own accord.
- The proof might use logic augmented with the axiom of induction to cover arbitrarily many cylinders to be stacked.
- That sort of proof of correctness using a "Fregean" formalism (in the sense of Sloman 1971) is what theoretical computer scientists mostly look for.
- However, it is unlikely that toddlers can use such forms of representation and reasoning: they turned up very late in human history, and normally turn up quite late in individual human development.
- Is there another way of thinking about the process that might allow a young child to grasp that it will necessarily achieve the goal, under certain conditions?

A more “intuitive” mode of reasoning?

Children (and most adults) apparently cannot construct logical proofs, but I conjecture that they **can** grasp the necessity of the outcome by using a different form of representation based on “analogical representations” (Sloman 1971, Peirce’s Icons?):



CONJECTURE:

various stages of the process are represented in a manner that is closer to how they are perceived visually than how they are described logically.

(‘Perceived visually’ has to be qualified for people blind from birth. But they still benefit from having had ancestors with vision if the blindness is caused by a peripheral defect not something central.)

The reasoning is easiest in the special case where the containers are first all arranged in a row, decreasing in width.

That configuration can be thought of schematically as having a widest cup on the left and a thinnest one on the right, with every adjacent pair (i,j) having the wider (i) on the left, but **without** representing specifically how many cups there are, nor how big any of them is, nor exactly where they are:

This requires a “generic” or “schematic” analogical representation – using “visual ellipsis”.

(Cf Kant vs. Berkeley, on generic triangle images.)

From that starting configuration (top image) you can envisage a first step of moving the widest cup to a new location (middle image) and then repeatedly taking the last cup on the left and putting it in the growing stack.

So there will be a generic situation that can be visualised, with the row of remaining cups, decreasing in width from left to right, and the already stacked cups somewhere separate – until one cup is left, and finally put inside the stack.

Representing a process visually

The crucial discovery is that there is an **invariant** structure while the repeated operations are performed when there are two or more containers left from the original row.

This can be represented by a **generic process step** that moves the leftmost cup from the row and puts it inside the last moved cup in the stack.

During this process there is still

- a row of remaining cups ordered by decreasing width from left to right
- Set of already stacked cups
- until there's only one cup left in the row, and putting that one in achieves the goal.

That's an invariant that can also be expressed logically, but probably not by a toddler!

I know some humans can do all this, but I don't yet know how to build AI systems with the appropriate forms of representation, self-monitoring capabilities, architecture, etc.

This process is similar in character to a proof in Euclidean geometry where a diagram (on paper or in the head) indicates the possibility of transforming one configuration to another, e.g. by adding construction lines, or by marking lines or angles equal, etc. (Sloman 1971)

Each such diagram, and each transformation, represents not **just one particular case** but **a whole class of cases**: the ontology of structures and processes is partly diagrammatically characterised. (How does that work?)

A more formal theory of container nesting

Some learners will develop mathematical ways of thinking, using mathematical formalisms and strategies, and thereby gain more general understanding of why something works.

For example, the learner may come to realise that the previously learnt successful strategies are **special cases** of more general strategies.

For example, stacking the nesting cylinders need not work monotonically from smallest to largest or from largest to smallest, because of the following facts about orderings of things and actions:

1. If n objects have different sizes and no two are the same size, then there is a decreasing-size-ordered sequence: $O_1, O_2, O_3, \dots, O_n$, where any two objects O_i, O_j , will be in one of two relations, $Bigger(O_i, O_j)$, or $Bigger(O_j, O_i)$, and each except O_1 has exactly one immediate predecessor and each except O_n has exactly one immediate successor: the sequence is “well-ordered” by size.
(Defining “immediate predecessor” and “immediate successor” is left as an exercise.)
2. The sequence exists independently of the physical arrangement of the objects.
3. For any object O_i , other than O_1 , its immediate predecessor can be inserted into it, as follows:
for every other object O_j
 - if $Bigger(O_j, O_i)$ then reject O_j (it's too big)
 - if O_k is found such that $Bigger(O_i, O_k) \& Bigger(O_k, O_j)$ then reject O_j (not immediate successor)
 - Otherwise insert O_j in O_i .
4. If $n-1$ actions are performed, where each action involves putting an object in its immediate predecessor (the next bigger), then the result will be a rearrangement of the objects so that each of the objects O_i , except for O_n , contains all the objects O_j such that $Bigger(O_i, O_j)$.

This formulation allows different orderings of the actions, but step 3 is slow and error prone.

Some informal experiments

The next few slides present a domain for **you** to think about, consisting of the configurations and transformations possible with an elastic band and a collection of pins.

This is one of many relatively small, restricted, domains in which one can interact with, play with, explore and learn from a subset of the environment.

Very young humans, and some other animals, over many months, interact with, play with, explore and learn from very much richer, more varied, more complex environments.

That process of conceptual, ontological, and “scientific” development is, I believe, hardly driven at all by biological physical needs for food, drink, warmth, physical comfort, etc. – though they have their role in other processes.

I have also deliberately abstracted from the role of **social** interaction during play since it is clear that in some species a great deal of exploration and learning can happen in periods of play **without** social interaction.

(I also think the role of imitation has been over-rated: you cannot imitate something done by another if you do not have the prior ability to do it yourself. If you are not so equipped, no amount of demonstration by another will teach you – e.g. to play a violin, place a puzzle piece in its recess, or to talk!)


I hope that your experience as an adult, playing in an accelerated and simplified way, with a partly familiar and partly unfamiliar domain of structures and processes involving pins and rubber bands, will help you to appreciate some of what happens during infancy and childhood – including some of the long forgotten things that happened during yours.

Try to be a child 1

A child can be given one or more rubber bands and a pile of pins, and asked to use the pins to hold the band in place to form a particular shape

Ideas for many games and explorations for children, using string, elastic, pins, stones, pencil, paper, etc. can be found in this excellent little book (Penguin Education series):

Jean Sauvy and Simonne Suavy,
The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology, 1974



-
- You can easily work out how many pins you need to form the rubber band into **a square** (The pins need not go through the elastic! Only sideways pressure is needed.)
 - How many pins will you need to make **a six pointed star**?
You can probably work out an answer by **reasoning** about stars, possibly imagining or drawing one.
A young child might find an **empirical** answer by exploring ways of pinning the rubber band into shape.
A younger child may be able only to discover the shape by chance or by copying someone else.
 - How many pins will you need to form **an outline capital “T”**?
 - What’s the **minimum** number required?
 - Are you sure? (A child may only be able to report results of actual trials.)
 - How can you be sure that you have found the minimum number required?
(Assume all the corners of the outline “T” are right angles (90°), ruling out a “T” made from three long thin triangles.)
 - Does the minimum number required change if the “T” is upside down? How do you know?


Try to be a child 2

Stretching a rubber band to form an outline capital “T” is easy.

Are any of the pins shown forming the “T” on the right redundant?

How do you know? Why don’t you have to try removing them to see what difference that would make? (Like a very young child?)

Can you see another way to make a capital “T” – using only five pins? [How?](#)



A capital “C” is possible, but not so easy.

How many pins would you need to make the outline of the “C” completely smooth: is it possible?


A child thinking about that might discover deep facts about continuity and curvature.

What’s the minimum number of pins required to form the “C”?

Is that a well-defined question?

Which of the outline capital letters can, and which cannot, be formed using pins and a single rubber band?

In particular, which of these: “A”, “B”, “D”, “E”, “F”, “G”, “H”, “O”,



A young child might have to do a lot of experimenting in order to answer these questions. You can probably answer them simply by thinking about them, visualising various possibilities. What goes on in a brain, or mind, doing that?

[How did you acquire such competences, and what forms of representation, mechanisms, and architectures make those competences possible?](#) As far as I know, nobody has answers.

Rubber band (toddler?) theorems

A child playing with shapes that can be made using a rubber band and pins (or pegs on a peg-board – a more constraining environment)

- may discover, by observing what happens, many ways of rearranging the configurations, and
- explore consequences of various combinations of arrangements and rearrangements
- and then make some new discoveries about the effects of those sequences.

For example, things to be learnt could include:

- There is an **area inside** the band and an **area outside** the band
- The possible effects of moving a pin that is **inside** the band towards or further away from other pins inside the band. **(The effects can depend on whether the band is already stretched.)**
- The possible effects of moving a pin that is **outside** the band towards or further away from other pins inside the band.
- The possible effects of adding a new pin, inside or outside the band, with or without pushing the band sideways with the pin first.
- The possible effects of removing a pin, from a position inside or outside the band.
- Patterns of motion/change that can occur and how they affect local and global shape (e.g. introducing a concavity or convexity, introducing or removing symmetry, increasing or decreasing the area enclosed).
- It's possible to cause the band to cross over itself. **(NB: Is an odd number of crossings possible?)**
- How adding a second, or third band can enrich the space of structures, processes and effects of processes.

You can play with a virtual rubber band here, http://nlvm.usu.edu/en/nav/frames_asid_172_g_2_t_3.html and here but only making convex shapes: http://nrich.maths.org/public/viewer.php?obj_id=4901

Fallibility of mathematical reasoning

It is well known that mathematicians can make mistakes of various kinds.

Imre Lakatos in **Proofs and Refutations** demonstrated this in connection with the complex history of Euler's theorem about polyhedra.

- He reached the conclusion that mathematics is “quasi-empirical”.
- In a separate slide presentation (in preparation) I shall produce more examples of types of “bugs” in reasoning that occur, and need to occur as part of the process of development of understanding of spatial structures and processes and other aspects of mathematics.
- The processes of fixing such bugs are partly analogous to the process of fixing bugs in designs for computer programs.
- In both cases, the fact that we can get things wrong, and find ways to fix them does not imply that what we are learning in the process is merely new contingent facts about the world.

There are deep problems here that bring together issues in

- philosophy of mathematics,
- philosophy of science,
- developmental psychology,
- Artificial Intelligence,
- computer science, and
- requirements for a satisfactory framework for teaching mathematics, based on a deep understanding of the processes of learning, which has so far eluded educators and psychologists.

Producing, detecting, fixing bugs in mathematical proofs

What is the minimum number of pins required to produce a star, using a single rubber band, where

a star is defined to be a polygonal figure with alternating convex and concave corners?

There is a plausible way of reasoning about the problem that leads to the conclusion that the minimum number of pins is six, producing a three pointed star, with three convex and three concave corners.


(Left as an exercise for the reader.)

In a discussion of the problem Zeyn Saygol produced the counter-example on the right, constructing a star using only five pins.

A number of options present themselves

- Reject the **original** proof as invalid, and its result as false, replacing it with a new conjecture that the minimum number of pins required is five (but see next slide).
- Reject the **new** proof by arguing that the final figure shown does not fit the Intended interpretation of the definition of a star, since it requires the boundary to cross over itself, in several places. The definition can then be made more precise by requiring a star to be a **simple** closed polygon.
- Make the original conjecture more precise, e.g. by constraining the allowed methods of using pins to create a new vertex, so as to rule out this construction.
- others ???

Compare the discussion of counter-examples in Imre Lakatos' **Proofs and Refutations**.



Another “star-theorem” counter-example


In a discussion prior to the AISB2010 symposium on mathematical cognition, after I described the problem of finding the minimum number of pins required to make a star, using a rubber band, Thomas Joyce pointed out that there is yet another counter-example to the claim that at least six pins are required, as shown on the right.

The diagram shows how to produce a three-pointed star, starting with a triangle formed using three pins, then adding a fourth pin in the middle and pulling each of the three sides over the central pin.

Question: does this meet the definition of a “star” as a [polygonal figure with alternating convex and concave corners](#)?

- Is it a polygonal figure?
- Does it have alternating convex and concave corners?
- Would it make any difference if the pins had infinitely small diameters?
- This figure cannot be constructed by adding pins and using them only to push or pull portions of the rubber band: here it seems to be necessary to lift at least two portions over the pin. (Why?).

Does that suggest that the statement of the theorem should say more about how the pins are used?



It's not just rubber bands

The questions presented above about what can be learnt by playing with rubber bands and pins are not unique.

Similar questions can be asked and answered as a result of playing with and exploring

- many **materials** of different kinds
(including string, wood, paper, water, sand, cloth, crayons....),
- **structures** involving those materials
(including marks on surfaces),
and
- **processes** in which structures and relationships change.

As adults, we can think explicitly about many of these processes, though we do not have access to the mechanisms and representations we use in doing that (though some people may **think** they do).

There are deep continuities and deep differences between adult processes and infant and toddler processes, and unanswered questions about the nature of the progression to adult competences.

Theories of kinds of stuff, bits of stuff, bits of process

Experiments with pins and rubber bands involve **interactions** between bits of stuff of different kinds – **the interactions are bits of process**.

The bits of stuff include:

- portions of the elastic band – all flexible and stretchable
- rigid, pointed, graspable pins, with a pointed end and a blunt end, and other portions
- the surface on which the band and pins lie, and into which the pins can be pushed
- the materials composing the experimenter's hands and fingers.

The bits of process involve changes described in previous slides, including both local changes (e.g. **pin pushes piece of rubber sideways**), and more global changes

(e.g. **a straight line becomes bent, or *vice versa*, the rubber becomes more taught and more resistant to further pushing, an asymmetric shape becomes symmetric, or *vice versa***).

Someone playing with pins and bands could acquire a moderately complex, but finite, “**micro-theory**” about many local structures and processes that can occur.

The theory can generate predictions about results of indefinitely many (but not all) different combinations of actions involving the rubber band and pins

- The predictions will use topological and geometrical reasoning, not just observed correlations.
- This is a form of intuitive child science, that might later be fully explained by adult official science (physics, materials science, plus mathematics).
- Compare the work by Pat Hayes on “The naive physics manifesto”.
- The theory will not predict exactly when the band will break if stretched a long way. (Why not?)

Alternatives to empirical knowledge

Empirical beliefs normally cannot be **proved** to be correct.

Not even in science, contrary to popular misconceptions. See the presentation: “What is science?”

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk18>

Contrast empirical information with:

- **Necessary truths** – some derivable using only logical or mathematical methods.

Examples – nothing could make these false:

P or (not P), not(P and (not P)), Containment is transitive, $3 + 5 = 8$

If P is true, and (P logically implies Q) is true, then Q is true.

If Q is false, and (P logically implies Q) is true, then P is false.

NB: necessity of the content does not imply that there are infallible modes of discovery. (See Lakatos).

- **Intermediate case (Semi-necessary/Relatively-necessary truths):**

Truths logically/mathematically derivable from the best available deep, general, well-established theory about how some important part of the world works. E.g.

Any portion of pure water contains twice as many hydrogen atoms as oxygen atoms.

Electric currents flowing through a metal produce heat.

A rubber-band star needs an even number of pins. (Or is that necessary?)

- **“Framework Theories”** (in one sense of the term)

Unavoidable/indispensable (implicit) theories required for perceiving, thinking, learning, acting, etc.

Explaining precisely what such a theory is would require us to build something like a working robot using one: too hard, at present!

So instead I'll merely offer sketchy examples, below.

Varieties of semi-necessary knowledge

Semi-necessary information includes deep, general, well-tested theories about how some important parts of the world work.

It also includes whatever is derivable from the theories, using derivations that depend on the form in which the information is encoded.

There are different types of semi-necessary knowledge, including the following:

- **Official-adult science**

Whatever is in, or is derivable from, the best established general theories about the nature of space, time, and the things that can occupy space and time.

That would include modern physics, chemistry, and biology, for example. (What else?)

This is the most explicitly formalised type of theory, and many of the derivations use logical and mathematical forms of inference.

- **Unofficial-adult science**

This includes a vast amount of information about kinds of things that can exist, kinds of process that can occur, and constraints on what can occur. (I.e. a rich ontology, much of it unarticulated.)

There is a core **shared** among all adults, and also much that **varies** according to culture, professions, and individual interests and histories: e.g. rubber-band science must be relatively new.

The contents of this sort of science are not usually articulated explicitly: much is represented internally implicitly in a usable form, and often used, though rarely communicated – e.g. craft knowledge.

Everything derivable (how?) from the contents of this science is also part of this unofficial science.

- **Baby-toddler science**

This includes precursors of adult-unofficial science, developed in early life – like rubber-band science.

Baby-toddler semi-necessary knowledge

Official-adult science has received most academic attention.

Baby-toddler science has only been investigated in narrow fields, in a shallow way, depending on changing fashions in psychology.

Unofficial adult science is somewhere in between (some of it studied in anthropology and ethnography).

A deep study of Baby-toddler science requires investigation of environment-driven requirements for the design of **working** animals, potentially testable in robots that develop as many animals do, e.g. through play and exploration in the environment.

Baby-toddler science includes deep and widely applicable theories developed in infancy and childhood that make sense of the environment, including:

- **theories about kinds of stuff** that can exist
 - e.g. various kinds of solids, fluids, sticky stuff, etc., with various kinds of flexibility, elasticity, etc.
 - Some of this will be mentioned later.
- **theories about kinds of relationship** that can hold between bits of stuff of various kinds
 - e.g. relationships between rigid fixed pins and elastic bands in various shapes.
- **theories about kinds of process** (changing properties and relationships) that can occur – e.g. what happens if a pin is removed in the capital “T”? What can happen if you pull a piece of string?

An incomplete presentation on the **ontology** developed in baby-toddler science is here:

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

DRAFT: Assembling bits of stuff and bits of process, in a baby robot’s world: A Kantian approach to robotics and developmental psychology. (Help needed extending this.)

Why Is Progress in AI/Robotics Slow?

Usual (wrong or incomplete) answers

- Answer 1: It's not really slow – look at all the new results

Visual recognition, visual tracking, 3-D reconstruction, DARPA challenges, robot soccer, impressive robots, e.g. BigDog but these are all very limited

e.g. doing but not thinking about doing:

real-time control but not reflection on what was and wasn't done or why, or what could have been done.

- Answer 2: Progress is slow because the problems are really really hard, and in order to solve them we need the following new approaches ...
 - Biologically inspired mechanisms (neural nets, evolutionary computation....)
 - Biologically inspired morphology for robots (emphasis on “embodiment”)
 - New architectures
 - Development of quantum computers, or some other new form of computation
 -

These answers miss the main point.

An alternative answer - Progress is slow because:

- We have not understood the problems
- Most of the problems come from the nature of the environment
Which has changed in the course of evolution:
There's not much scope for intelligence if you are a micro-organism living in an amorphous chemical soup
- We need to understand the cognitive problems and opportunities provided by the environment.
 - These arise from the fact that the environment is structured in many ways – some of which, at least, can be discovered and used.
 - This requires something deeper than using statistics to learn correlations and probabilities (e.g. Bayesian learning)
 - Some species, though not all, can do this.
 - Humans go further, and discover what they have learnt, formalise it, and transmit it across generations.
 - And they use it to do increasingly complex things with little intellectual effort:
support for productive laziness is a key feature of mathematical competences

All organisms are information-processors but the information to be processed has changed and so have the means



Types of environment with different information-processing requirements

- Chemical soup
- Soup with detectable gradients
- Soup plus some stable structures (places with good stuff, bad stuff, obstacles, supports, shelters)
- Things that have to be manipulated to be eaten (e.g. disassembled)
- Controllable manipulators
- Things that try to eat you
- Food that tries to escape
- Mates with preferences
- Competitors for food and mates
- Collaborators that need, or can supply, information.
- and so on

Structure in the world: Challenges and opportunities

Wittgenstein: The World is the totality of things, not of facts (Tractatus)

We can do better!

The world is the totality of processes: actual and possible

- Processes of varying complexity are composed of recombinable process fragments
- Possible process fragments and ways they are combined change over time
 - bits of stuff of different kinds
 - bits of surface
 - various kinds of juxtaposition and composition – spatial and temporal
 - changes in various structural and other properties and relationships
 - * on different spatial scales (very small to very large)
 - * on different temporal scales (very fast to very slow)
 - various constraints on what does and does not change: **invariants**

As more and more fragments are composed to form larger and larger wholes, the variety of possibilities becomes greater and greater – exponentially, or worse.

How can an animal be ready for the next situation ...

The core challenge to be met

The variety of ways of composing kinds of fragments into larger (spatially extended, temporally extended) wholes produces **explosive diversity of structures and processes**.

- So learning isn't enough for coping with something unexpected, something beyond the envelope of previously encountered cases.
- An alternative is to develop ways to **work out** consequences of novel juxtapositions of structures and processes,
e.g. new assemblages of levers, gears, wheels, hooks, ropes and pulleys
- But this alternative requires expensive, bulky, energy consuming brain mechanisms (though birds seem to have found a more efficient form of implementation than mammals – and what about octopus brains?)
- Two biological coping strategies for a species:
 - (a) produce large numbers of cheap offspring – most of which fail to reproduce
 - (b) produce few, expensive, highly capable offspring, more of which reproduce
 - EITHER with mostly genetically fixed competences and gradual adaptability
 - OR with genetically provided meta-competences with rapid adaptabilityThe latter can be the basis of mathematical development.

Embodiment is not the solution: it defines the problem

There are two fashions related to emphasis on embodiment (in psychology, neuroscience, philosophy, AI, robotics:

- Try to copy biological **morphologies**, including sensors and effectors: copy interfaces
- Try to copy biological **computing mechanisms** (e.g. neural mechanisms).

CONJECTURE: This misses the point!

There has been convergent evolution towards virtual machines required for gaining information about and interacting with a complex and changing 3-D environment – but using **different body structures and brain mechanisms**, e.g. in

- mammal intelligence
- bird intelligence (Show crow movie?)
- octopus intelligence

But there is something only humans have, though some other animals appear to have parts/aspects of it: the ability to use interaction with the environment to develop proto-mathematical competences:

(a) abilities to perceive, understand and reason about novel structures, situations and processes I focus on this, today. (b) and (c) depend on it.

(b) The ability to transmit what has been learnt

(c) the ability to understand and incorporate what has been transmitted

Something most researchers have not noticed

There are hundreds, perhaps thousands, of examples of things learnt empirically in infancy and early childhood that can be re-conceived as topological, geometrical, or mathematical truths, or as derivable from a deep theory about the environment.

Different children discover different things, and the sequence can be very different from child to child.

Moreover

- The ability to achieve that transformation of knowledge
 - from **purely empirical**
 - to **necessarily true, or derivable from a general theory**has great biological significance:
 - for animals that need to cope with complex novel situations.
- There are some special cases where the theory developed is so deep and general that it cannot be refuted empirically: it is more mathematical than empirical:
 - e.g. an empirical discovery turning out to be a theorem in topology or number theory.
- Developing the ability to do this requires structural changes in the information-processing architecture: it is not there from birth.
- Yet the potential for that development comes from the genome.
 - This may be related to what Annette Karmiloff-Smith calls “Representational Redescription”.
(In *Beyond Modularity*)

Confusions about the importance of embodiment

Much recent research in AI/Psychology/Neuroscience/Animal Cognition and Philosophy emphasises “embodiment” and dynamical interaction between agent and environment.

Attention often focuses on how an individual’s specific sensory and motor modalities and body morphology relate to requirements for perception, control and learning during acting.

Emphasis on real-time architectures and ‘online’ control ignores abilities to:

- (a) perceive, think about, learn about, predict, plan, reason about things that are possible, but not happening, or are out of sensory range, or may have happened in the past (‘offline’ processing);
- (b) integrate information on various spatial and temporal scales, forming a spatially and temporally extended “episodic memory” of a portion of the environment, including places and things out of view.

The environment drove evolution of information processing mechanisms supporting (a) and (b) using amodal abstract representations of actual and possible structures and processes divorced from specifics of any morphology or set of sensors or actuators.

Conjecture: more advanced forms of these amodal mechanisms require layered information-processing architectures alongside and interacting with the evolutionarily older multi-modal sub-architectures for online control.

Some species have genetically **pre-programmed** layers, whereas others require higher level layers to **learn** to explore, plan, test hypotheses, revise strategies, etc..

This needs internal self-monitoring and self-control capabilities that include abilities to create new internal forms of representation and new theories, as well as new strategies – and support for **productive laziness**.

Relevance to Philosophy of Mathematics

IF

- we can demonstrate how mathematical competences can grow naturally
- out of biological requirements for interacting in a very flexible way with a wide variety of complex environments,
- using mechanisms and forms of representation that exist in pre-verbal children and some other animals,

THEN

- that will provide a novel contribution to philosophy of mathematics, including a vindication of Kant.

But the issues are subtle and complex,

e.g. because the switch in understanding from empirical to mathematical (non-empirical) is error prone, and “proofs” can be buggy, as Lakatos showed clearly in his discussion of the history of Euler’s theorem about polyhedra.

[See also the racing car example below.](#)

Although I can specify some requirements I cannot yet specify a full working design.

[Offers of collaboration welcome.](#)

A new cross-disciplinary collaboration

We can combine philosophy, robotics and biology.

In order to understand what mechanisms explain these processes, we need to understand what the biological requirements were that favoured evolution of competences that provide the basis:

- for being able to discover that a subset of things learnt empirically can be re-interpreted as necessary truths rather than contingent and empirical truths (either necessary relative to a “framework theory” or necessary relative to a powerful empirical theory).
- for being able to describe and use the necessary truths and learn some of them from other individuals instead of re-discovering them
- for **doing** mathematics by systematically organising the discovered truths.

Finding the requirements is a step towards creating working models that meet the requirements, and potentially explain the phenomena.

More on the learning process

- The stages of development at first **implicitly presume**, then later **explicitly articulate**, theories about processes and structures, and later still organise them into a system. (Not all individuals do all of this.)
- Some aspects of those theories will be empirical and potentially refutable (such as theories about what rigid, solid material objects do not do if left alone, e.g. split or merge) and others may be intrinsic to any process of learning about what can and cannot happen in a 3-D world where things can move.
- I call the latter “framework theories”: A long term research goal is teasing them out, and separating them from the empirical theories. (More on this later.)
- **NB: the processes are not guaranteed to be bug-free: errors can occur, but can be detected and fixed**

Hooks and Rings: The toddler railway mystery



At about 18 months this child is able to get a toy truck in his hand rotated so that the blank end is replaced by the end with a ring.
(Compare birds rotating a fish after catching it.)

But he is mystified when trying to join two **rings**, not noticing the **hook** on the truck in his right hand.

Later, he somehow worked out how to join the train parts.

Nobody saw how the change happened.

What might have changed in his brain?

(Or in a virtual machine running on his brain?)

Video here:

<http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/vid/>

Epistemic features of micro-worlds

In addition to positive and negative affordances for **actions** of various kinds, the environment also provides positive and negative **epistemic affordances**, i.e. :

information about the environment may be available, or obstructed, in the environment (e.g. in the optic array, as J.J. Gibson noted).

Learning about how epistemic affordances change is just as important as learning how physical configurations change.

Processes occurring in the environment can alter both types of affordance:

- If something is too far to reach with your hand, you may be able to crawl forward: in that case crawling forward can make the object reachable.
- If not all of an object A can be seen because another object B, which is convex, partly occludes it, then moving object B can cause either less or more of A to be seen, depending on the direction of motion.
- Likewise the perceiver can make more or less of A visible by moving left or right (perpendicular to a plane through the visible edge of A and the occluding edge of B).
- Attempting to push a rod through a hole in a wall can provide information about the hole's diameter.

Those generalisations may first be learnt by abstracting from experienced examples.

Initially they have the status of empirical knowledge: new examples could refute them.

Later some can acquire a new status because they are derivable from very general topological and geometrical features of space – and features of how information flows. (J.J.Gibson)


DEEP problems: How are such “laws” represented in a child's, or chimp's mind, or brain?
What changes when their status changes? (See the cup nesting example.)

Getting information about the world from the world

Toddler theorems:

Things you probably know – and children can learn:

- You can get new information about the contents of a room from outside an open doorway
 - (a) if you move closer to the doorway,
 - (b) if you keep your distance but move sideways.Why do those procedures work? How do they differ?
(Theory: visual information travels in straight lines.)
- Why do perceived aspect-ratios of visible objects change as you change your viewpoint?
 - A circle becomes an ellipse, with changing ratio of lengths of major/minor axes.
 - A rectangle can become a parallelogram or a trapezium.
- Why do you see different parts of an object as you move round it?
- When can you avoid bumping into the **left** doorpost while going through a doorway by aiming further to the **right** – and what problem does that raise?



Non-epistemic affordances concerned with opening and shutting micro-worlds:

- In order to shut a door, why do you sometimes need to push it, sometimes to pull it?
- Why do you need a handle to pull the door shut, but not to push it shut?
- How you could use the lid of one coffee tin to open the lid of another which you cannot prise out using your fingers?

Developing meta-knowledge: mathematician style

A child, or a mathematician, can approach these problems in different ways, from which different things can be learnt.

- It is possible to explore possibilities by physically changing the environment – moving pins and elastic bands around, or tracing routes through the bridges with a pencil.
- It is possible to do the same explorations **without** acting on the world, **merely** thinking about the actions (Sloman 1971) – though in some cases it is hard to think about effects of complex sequences without actually doing them, because of human short term memory limitations (not shared by computers! – though computers cannot yet do much spatial thinking: only low-level simulating.)
- It is possible to do unsystematic searching through possible solutions to a problem in the hope of finding a solution: e.g. moving pins to different places, starting routes in different places and following different branches.
 - Finding a solution through unsystematic search proves that one exists.
 - Failing to find a solution through **unsystematic** search proves nothing.
- It is possible to notice that unsystematic searching can lead to wasted efforts because some combinations are tested more than once.
- Unsystematic searching can also lead to some combinations of actions not being tried at all.
- It is sometimes possible to devise a way of organising a systematic search to make sure that all combinations are tried and no combinations are omitted (this usually requires making use of **orderings**.)
 - Failing to find a solution through **systematic** search proves that there is no solution.
- It is sometimes possible to discover that both systematic and unsystematic searches are wasteful because there is a way of thinking about the problem that reveals in advance what the outcome of a search must be, e.g. proving that certain problems **cannot** be solved.

That is what Euler did – and mathematicians generally try to do. **How might robots do this?**

Axioms for rubber banditry: Adult & Child theories

An adult mathematician or logician could probably, in a few hours, or maybe a few days, develop an axiomatic theory of the rubber-band plus pins (RB+P) domain.

Adult mathematician

- There would be a choice between a fully metrical theory (lengths, angles, areas, etc. have numerical values) and a semi-metrical theory, lengths, angles, areas can be partially ordered only.
- The theory might get a lot more complex if more than one rubber-band was allowed.
- Perhaps using a subset of (e.g.) Hilbert's axioms for Euclidean geometry, plus axioms for **processes** altering geometric structures, e.g. bending lines, stretching lines, rotating lines, etc.), minus axioms implying infinite divisibility or infinite extension. http://en.wikipedia.org/wiki/Hilbert's_axioms
- Dynamics of stretched elastic might need several more axioms.

Child mathematician

- My guess, without having counted, is that a child might learn between fifty and about two-hundred (initially empirical) facts about the RB+P domain, many of which would be common to other domains, some of which may be false generalisations based on incomplete experiments, etc.
- But the child would not encode these facts in an axiomatic predicate calculus formalism.
- I don't know what alternative formalism is available to a child, but it could include a (highly context-sensitive) formalism for expressing how to **do** things (perhaps partly shared with apes?), and could make use of something like abstract pictures or maps for reasoning with, as suggested below.
(Cf. Sloman 1971. Jean Mandler, among others, also claims something like this.)
- Some children would discover **dynamic** properties – e.g. use of stretched rubber as catapult!
- **How** a child or any other animal does any of this (e.g. what their brains do) remains unexplained.

Try some other micro-worlds

The domains presented so far are just a small sample of micro-worlds to be found embedded in everyday situations.

Other domains that a child, or pet animal, or domestic robot might play with and explore:

- Sand, in a box, with or without a bucket and spade
- Sand on a beach, with or without waves coming up periodically
- Lumps of plasticine, all the same colour.
- Lumps of plasticine of different colours.
- A collection of stones of varying shapes and sizes.
- A collection of wooden cubes, all the same size.
- A collection of wooden cubes, of different sizes.
- A collection of rectangular blocks of different shapes and sizes.
- All the above with 'boxes' instead of 'cubes' and 'blocks'.
- A plateful of food
- Water, in containers on the floor. Try adding cotton-wool? tissues? ping-pong balls? marbles?
- Water, in which you are immersed, in a bath.
- A baby lego set.
- A meccano set.
- A doll with changeable wardrobe.
- Paper dolls with changeable paper clothing. <http://marilee.us/paperdolls.html>
- Sheets of paper, with or without scissors, pencil, a ruler, glue,
- Sheets of tissue paper.
- Clothing, eating utensils, furniture, human body-parts, playground entities, ...

Toddler reasoning quirks: reasoning about potentially colliding cars

A child can learn to use a mixture of simulation and symbolic/verbal reasoning to draw conclusions and to justify conclusions.

But that learning takes time and the process can go through buggy stages as illustrated by this example.



The two vehicles start moving towards each other at the same time.

The racing car on the left moves much faster than the truck on the right.

Whereabouts will they meet – more to the left or to the right, or in the middle?

Where do you think a five year old will say they meet?

One five year old's spatial reasoning



The two vehicles start moving towards each other at the same time.

The racing car on the left moves much faster than the truck on the right.

Whereabouts will they meet – more to the left or to the right, or in the middle?

One five year old answered by pointing to a location near 'b'

Me: Why?

Child: It's going faster so it will get there sooner.

What produces this answer:

- Missing knowledge?
- Inappropriate representations?
- Missing information-processing procedures?
- An inadequate information-processing architecture?
- Inappropriate control mechanisms in the architecture?
- A buggy mechanism for simulating objects moving at different speeds?

Partly integrated competences in a five year old

The strange answer to the racing car question can perhaps be explained on the assumption that the child had acquired some competences but had not yet learnt the constraints on their combination.

- If two objects in a race start moving at the same time to the same target, the faster one will get there first
- Arriving earlier implies travelling for a shorter time.
- The shorter the time of travel, the shorter the distance traversed
- So the racing car will travel a shorter distance!

The first premiss is a buggy generalisation: it does not allow for different kinds of ‘race’.

The others have conditions of applicability that need to be checked.

Perhaps the child had not taken in the fact that the problem required the racing car and the truck to be travelling for the same length of time, or had not remembered to make use of that information.

Perhaps the child had the information (as could be tested by probing), but lacked the information-processing architecture required to make full and consistent use of it, and to control the derivation of consequences properly.

Is Vygotsky’s work relevant? Some parts of Piaget’s theory of “formal operations?”

The (huge) micro-world of number competences

It is often wrongly assumed that the ability to perceive differences between sets of objects of different numbers (“subitizing”) involves understanding what numbers are.

However that is only a **perceptual** ability which has a loose connection with numbers.

Even showing that a child can learn associations between a sound triplet and a visual or haptic triad does not show that the child is doing any more than associating percepts in different modalities.

It need not involve any understanding of what can be **done** with numbers.

Understanding numbers (the natural numbers) requires at least:

- the ability to generate repetitive sequences, individually and in parallel, and to synchronise them
- the ability to generate, make use of, and reason about one-to-one mappings
- the ability to generate a learnt (initially finite) sequence of unique names/noises/labels, e.g. “one”, “two”, “three”, or “I”, “II”, “III”, etc.
- the ability to set up one-to-one mappings between arbitrary entities (e.g. stones, apples, steps or other actions, revolutions of a wheel, flashes), and an initial sequence of the number names
- the ability to use such mappings, **in combination with different stopping conditions** to perform various tasks (e.g. fetch enough spoons for everyone at the table, fetch three apples) and answer questions (how many people are there?)
- the potential to replace the finite names with an infinite generative scheme
- and potential to add a host of other competences that involve increasing understanding of numbers

See Chapter 8 of *The Computer Revolution in Philosophy* (Sloman, 1978)

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

Invariance of counting results

A child who has learnt to use those competences to answer “How many” questions may discover empirically that

counting the same set of objects (e.g. fingers on a hand) from left to right gives the same result as counting that set of objects from right to left.

Initially it may merely be an interesting empirical discovery.

Later, the child learns that the irrelevance of counting order to the result of counting an unchanging collection of objects is a necessary truth. How?

Can you prove that any two orders of counting the same fixed set, without repetitions or omissions must end with the same number?

It requires understanding how to set up a one-to-one correspondence between the elements in the two one-to-one correspondences: i.e. between two counting processes.

Detecting and explaining the necessity requires a more sophisticated information processing architecture than merely being able to perform the tasks listed above, and discovering empirically that order of counting does not matter.

I suspect that architecture also develops later, as part of general development of meta-semantic competence and meta-management (self-monitoring and self-control of many sorts).

All this concerns the seeds of an understanding of various kinds of infinity, discussed briefly later.

Chapter 8 of Sloman 1978 also showed how a child can discover an algorithm for answering questions of the form ‘What comes before X?’ Alternatively, the answer once found can be stored, associated with X.

Something non-empirical can be discovered empirically

Many things are discovered empirically that can be shown later to have a non-empirical proof, e.g. using logic, mathematics, or conceptual analysis to reach the same conclusion.

- It is possible to discover **empirically** that a logically true proposition, e.g. of the form **P or (not (Q and not Q) and not P)**, is true, by establishing P empirically, or by establishing Q and not-P empirically.
- A child playing with blocks by laying them out in various patterns, may discover empirically that a collection of 6, 8, 15 or 20 blocks can be arranged in a rectangular array of rows and columns, whereas a collection of 5, 7, 11 or 17 blocks cannot by exhaustively trying all the possible regular arrangements and always finding one or more surplus blocks – Later the child may learn about prime numbers.
- Someone could conduct a survey asking individuals
 - which people they respect
 - whose assertions they are likely to believe

And then discover using statistical analysis of the data:

People are more likely to believe new information from people they respect.

Later the researcher could learn to use conceptual analysis instead!

For more on this see

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#701>
'Necessary', 'a priori' and 'analytic', *Analysis* 1965.

The idea of a “framework theory”

The attempt to generate explanations unifying things that have been discovered empirically can lead to an explanatory theory (often referring to unobservables).

Where it clearly makes sense to formulate rival theories that could be investigated empirically – the theories themselves are empirical.

So their consequences have only a **relative** kind of necessity.

We could call them “semi-necessary”.

But when there seems to be only one possible theory – no alternative is possible – that could be called **a framework theory**.

We can be mistaken about whether something is a framework theory.

Compare Euclidean geometry: the parallel axiom seems to be dispensable.

But what about Euclidean geometry without the parallel axiom?

Compare topology, set theory, logic....

CONJECTURE: There are framework theories required for a robot learning about our physical environment.

Are there framework theories for thinking about thinking, cognition, etc.

(in oneself, or others)?

What could be in infant/toddler “framework theories”?

Here are some first draft speculations about contents for infant/toddler “framework theories”, though possibly not at the very earliest stages:

- The world contains various 3-D volumes occupied by **kinds of stuff**, with various properties and relationships, and some volumes that are empty.
- Portions of those kinds of stuff (i.e. bits of stuff) occupy space, have size, shape and orientation, and endure for varying time intervals.
- Over time they may be involved in various types of events and processes, in which their properties, relationships, parts, locations, orientations, can change.
- A changing subset of what exists and of what happens can be perceived.
- As suggested by Kant, an infant learner needs an implicit “framework theory” about the topology of space and time, some of their metrical structure, and assumptions about the role of causation in the world. (Exactly what is needed remains to be specified.)

Example theoretical assumptions:

containment (spatial or temporal), **being longer than** and **having a one-to-one correspondence** are all **transitive** – i.e. they are all instances of **$R(a,b) \ \& \ R(b,c) \ \rightarrow \ R(a,c)$** .

Bits of stuff, events, processes, portions of space, portions of time, and many abstract entities can be involved in one-to-one mappings (crucial for both mathematics and music).

- Some of the details of the above forms can be learnt by observing & experimenting.
- General laws applicable to the contents of the world can be discovered:
they cannot be discovered by any algorithmic process, but it is possible to make guesses and test them, weeding out bad theories.

See also John McCarthy’s “Well designed Child”, mentioned above.

Things apparently needed in framework theories

Here are some examples.

- There is a portion of space linking every two disjoint portions of space.
- Every two time intervals are inside at least one longer time interval that includes them.
- Every portion of space (1-D, 2-D, or 3-D) is part of a larger portion.
- Every interval of time is part of a larger interval.
- Every non-overlapping set of time intervals is totally ordered.
- The orderings of any two non-overlapping sets of time intervals are consistent.
- If a bit of stuff S exists throughout all of an ordered, overlapping, sequence of time intervals T_1, T_2, \dots, T_n , then there is at least one ordered and connected sequence of spatial regions occupied by that bit of stuff such that at each T_i it is in one of those regions.
- Different bits of stuff can start in the same region at the same time and one of them get to the last in the connected sequence of regions at an earlier time than the other. (Different speeds of travel.)
- Some bits of stuff can share regions with others, i.e. allow interpenetration (e.g. liquids) whereas others cannot?
- The bit of space occupied by anything at any time has a shape.

How do toddlers, reason?

Also crows, chimps, hunting mammals, etc.,

CONJECTURE They use non-verbal, non-logical languages:

Before languages for communication evolved, “internal languages” evolved for

- contents of perception
- representation of goals, desires, plans etc.
- representation of reusable generalisations
- episodic memories

These forms of representation can be called “languages” because they support:

- structural variability
- (context sensitive) compositional semantics
- manipulability extraction, recombination of parts, modifying ...
- addressability and matching
- coping with novelty (generativity)
- and therefore: reasoning

Such things seem to be available to non-verbal animals and to pre-linguistic humans.

For more on this see my 1971 IJCAI paper (and Chapter 7 of *The computer revolution in philosophy* (1978))

<http://www.cs.bham.ac.uk/research/projects/cogaff/04.html#200407>

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

Compare Jean Mandler’s work. (Ref)

NB: as noted in Sloman 1971: spatial reasoning can use [external](#) as well as internal [diagrams](#).

Some wrong assumptions about mathematics

- Being able to perceptually distinguish patterns of different numerosity (subitizing) shows some understanding of numbers: NO

The ability to subitize is neither necessary nor sufficient for understanding numbers

Understanding natural numbers (ordinal and cardinal) is based not on perceptual patterns but on understanding one-to-one correlations, e.g. between steps in synchronized sequences, or between collections of physical or abstract entities.

Compare Rips et al. in BBS 2008

- Mathematical competences are all numerical competences: NO

Some are competences relating to geometry, topology, logic, mechanics, plan formation, problem solving, and more.

Some high level conjectures

These need to be tested both empirically and by building working models (very hard to specify precisely).

- Important forms of learning, in human infants/toddlers, and possibly some other “altricial” species occur in genetically “staggered” phases
- An early phase involves doing various things in the environment and finding out which actions and observed processes produce or prevent various sorts of happenings.

The results of this learning mainly take the form of action competences, and prediction competences: they are not necessarily stored in a directly accessible propositional form that can be used for reasoning in any context.

- When enough has been learnt to allow various specific goals to be achieved a new form of learning starts which plays with both internal state changes and and external actions, forming conjectures, testing them, debugging them.
- This process can lead to the development of theories about the environment.

The results of this learning can, eventually, lead to information stored in forms that can be used in a variety of contexts, including, at a later stage, communication as well as reasoning.

This process may depend on social facilitation (scaffolding).

- At a still later stage the information can be more systematically organised in a deductive form.

This also depends on social facilitation.

- There are individual variations in the processes (e.g. autism)

Intelligence as Productive Laziness

As Alan Turing pointed out in 1950 “intelligence” is too vague a notion to be worth trying to define.

Nevertheless we can observe that an important aspect of intelligence is **productive laziness**.

Productive laziness involves finding ways to get things done with minimal effort (or less effort than previously).

Compared with mechanisms that explore and use readily available alternatives, such as exhaustive or random search for a solution to a problem, mechanisms required for productive laziness are more complex, performing computations of greater abstraction and sophistication.

They must have evolved later, and in far fewer species (near the top of a food pyramid).

In mathematicians, productive laziness reaches extremes: a mathematician can, in a few short steps, answer questions about infinite sets, e.g. **Is there a largest prime number?**

To see some simple and not so simple ways to prove that the answer is No, look at

<http://primes.utm.edu/notes/proofs/infinite/>

Could you answer the question about pins needed for a six pointed star? You did not search through the infinitely many ways of shaping a rubber band into a six-pointed star.

So you too are a productively lazy mathematician, covering infinitely many cases in a few thought steps.

What goes on in brains doing such things?

Or in minds – the virtual machines running on brains?

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#wpe08>

Main theses – aiming to unify a lot of empirical data

- Humans (+perhaps some other species) acquire various kinds of **empirical** information (e.g. by playing) then LATER develop a new **derivation** of the information, removing or changing its status to non-empirical – having a kind of necessity (defined later).

They may not realise they are doing this, unless they study mathematics – perhaps not even then!

- The forms of representation, perceptual and other mechanisms, ontologies, and information-processing architectures that make the transformation possible, evolved to meet biological requirements imposed by complex and changing 3-D environments.

I'll try to show how this is important for the ability to produce creative solutions to novel problems, without having to do inherently slow statistical learning or dangerous testing where errors can kill.

- The competences required for this are not all present at birth: they develop in layers driven partly by the environment. See Chappell & Sloman

Natural and artificial meta-configured altricial information-processing systems, IJUC 2007

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>

- In humans, those biological competences provide the basis of the ability to do mathematics, and some of that ability exists (unrecognised) even in young children.

I'll introduce the notion of a “toddler theorem” and give examples.

- Understanding the biological origins of mathematical competences provides support for Immanuel Kant's philosophy of mathematics,

wrongly thought to have been refuted by the discovery that physical space is non-Euclidean.

- **I hope the ideas can be tested and demonstrated one day in a “baby” robot.**

- I welcome collaborative research in this area. Any developmental psychologists interested?

Two very different views of mathematics

Two intellectual giants of the past century differ strongly:

Bertrand Russell (philosopher and logician): Mathematics may be defined as the subject in which we never know what we are talking about, nor whether what we are saying is true. (In *Mysticism and Logic* 1918)

Richard Feynman (Theoretical physicist): To those who do not know Mathematics it is difficult to get across a real feeling as to the beauty, the deepest beauty of nature. ... If you want to learn about nature, to appreciate nature, it is necessary to understand the language that she speaks in. (In *The Character of Physical Law*)

My claim: developmental psychologists, roboticists, and philosophers of mathematics need to understand how Feynman's views relate to child (and future robot) development.

A requirement for neuroscience is to explain how brain mechanisms make this possible.

Can anyone explain how a genome encodes the mechanisms that build new layered mechanisms driven by interactions with the environment?

Differences between Russell and Feynman are discussed in a little more detail with more examples, in <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#math-robot>

These quotes about mathematics, many other quotes, and lots of mathematics-related fun and information, can be found in:

<http://www.cut-the-knot.org>

Russell, Feynman and Kant

Russell (a modern logicist Humean) viewed mathematics as **the investigation of implications that are valid in virtue of their logical form, independently of any non-logical subject matter.**

Feynman, a leading physicist, viewed mathematics as **“the language nature speaks in”**.

Kant's view, that mathematical knowledge is synthetic and non-empirical (but not innate), can be interpreted as a modified version of Feynman's view, something like this:

The ability that develops in a human child (and possibly in some other species), namely to see, manipulate, predict, plan and explain structures and processes in nature, requires a collection of information-processing mechanisms that turn out also to be the basis for the ability to make mathematical discoveries, and to use them.

(Closely related to the ability to notice and investigate philosophical problems: a topic for another day.)

CONJECTURE

This process could be replicated in a suitably designed robot.

Doing that could help to substantiate Kant's view (as well as contributing to developmental psychology).

If it turned out that all of a robot's knowledge about the world had to be learnt/statistical/probabilistic/correlational, then that could be a refutation of Kant's ideas (as I understand them).

The next few slides introduce some technical ideas about kinds of information, informally.

Why the designer stance?

It is important to stress that an adequate theory should provide a design for something that **works**?

A deep study of Baby-toddler science requires investigation of environment-driven requirements for the design of **working** animals, potentially testable in robots that develop as many animals do, e.g. through play and exploration in the environment.

This is important because researchers often put forward what are intended as explanatory theories, but are actually nothing more than empirical correlations (even well confirmed empirical correlations), which do not explain anything.

Suppose it was found that children who had played with certain toys, and who had no known physical disabilities or brain abnormalities, could learn to produce proofs in geometry much faster than others.

That would **not** explain how anyone produces a proof, or why playing with those toys accelerated learning.

On the other hand

- if you produce a design for a machine that can play with the toys, and
- demonstrate how such play alters internal information structures, algorithms, and perhaps architectures, and
- show how those alterations improve the rate of mathematical problem solving in such a machine,

then you would have a deep explanatory theory which could, for example, indicate ways of improving the design of the toys or the manner in which they are introduced.

It could also inspire research into how those processes work in brains.

What about disabled learners?

Which of my claims about what can be learnt by a normal human child, about structures and processes that can occur in the environment, would have to be changed for a child born blind, or without arms, or with defective haptic sensors, or with cerebral palsy preventing control of hands?

In particular, is it **just the environment** that provides the sources of information to trigger the kind of cognitive development outlined here, or is the possession of **a specific sort of body** required?

This is an empirical question that could be partly answered on the basis of literature on achievements of blind mathematicians, blind musicians, thalidomide babies, people born with missing or truncated limbs (e.g. the artist Alison Lapper), and others.

Conversations with clinical developmental psychologist Gill Harris, have taught me that there is much more in common between the cognitive and other competences that can be achieved by normal and disabled children of various kinds than is generally known.

The time required, the routes to the competences, and the amount of help required from others will, of course, vary much more than the end results.

The children with missing sensors or missing or defective limbs, etc., nevertheless have **brains** that evolved in ancestors well equipped to learn about a rich and complex 3-D environment.

For a more detailed discussion see

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804>

Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress (To be published in 2009).

Can framework theories develop?

An implicit framework theory (about a space of possible environments, and ways of learning about them) seems to be a prerequisite for any learning at all to happen.

However we must distinguish two possibilities, the first often taken for granted, but probably wrong, the second much more complex and unobvious:

1. Unchanging correct framework:

The initial (innate) framework theory is correct and is retained through all subsequent developments.

2. Buggy developing framework:

Initial theories are incorrect and incomplete, but useful for launching and steering the learning process.

- The learning process includes debugging and replacing parts of the initial theory.
- In each learner, an improved, more general, less buggy, framework theory develops stepwise out of the interactions between the initial incorrect theory and the environment. (See epigenesis below.)
- In some cases effects of cultural evolution and development can play a part in the process.

In case (2) the genetically produced initial theory is wrong, but was selected by evolution as an effective starting point for development towards a better framework theory, enabling various kinds of empirical theories to be based on play, exploration, observation,

How the changes are made is a topic still to be investigated: probably a sequence of processes of bootstrapping, testing, debugging, revising, extending,

A “toy” example – Sussman’s HACKER program: *A computational model of skill acquisition*, Elsevier, 1975 (More recent experiments use evolutionary algorithms.)

Framework theories for infants, toddlers and adults

As Kant noted, experience without any prior information is impossible:

- (a) spatial perception requires the ability to organise information spatially, at least in 2-D, and possibly 3-D;
- (b) temporal experience requires a grasp of temporal relations and spatial **changes**;
- (c) learning about causation presupposes some prior concept of causation.

It is very hard to tell what information is being acquired, derived, stored, recombined, matched, or used by newborn animals.

Eye-movement recordings give only very shallow information, and require grossly implausible assumptions (e.g. a neonate thinks a bit like an adult, has expectations, and can be surprised).

Neisser and McCarthy suggest we can learn what information processing might occur in very young children by studying the environment they evolved to interact with, and trying to understand its information-processing demands, including what information needs to be represented and how it needs to be manipulated and transformed in order to perform various tasks.

Unfortunately, experience shows that it is hard to do that: many researchers don't know what to look for in the environment that can produced demands on information processing mechanisms: "ontological blindness"

Information thought to be empirical can change status

What starts off as apparently purely empirical can turn out to be semi-necessary, e.g.

- an observed correlation turns out to be derivable (logically, mathematically) from a deep, widely applicable theory that explains a wide range of phenomena better than any other theory (e.g. current atomic theory of matter, or toddler physics).

What starts off as apparently purely empirical can turn out to be non-empirical, e.g.

- provable by logic or mathematics
- part of, or derivable from, a “framework theory” – something that is required for everything else.

Examples were given in the rubber-band domain, bridges domain, and other domains mentioned above.

That such transitions can occur is an important, largely unacknowledged, fact about development in humans – and maybe some other species?

I think this fact is in part what underlies Immanuel Kant’s philosophy of mathematics.

It seems to have been noticed by Jean Piaget, but not described accurately.

I shall give some examples, then return to a more general discussion.

I don’t yet have a working model of how a child develops so as to discover a new status for old information, only some hunches....

But we can specify some **requirements** for mechanisms capable of supporting what happens, and also explain why this has biological utility – partly explaining why such mechanisms were produced by evolution.

How much of this occurs in other species is unknown.

How can we test what a child, or robot thinks?

Very young children are inscrutable: what they can communicate is grossly impoverished compared with the information processing they must be capable of – including building an architecture.

Involuntary behavioural indications of contents of perception, learning, motivation, deciding, reasoning, planning will be seriously impoverished: behaviour at most reflects aspects of final stages of the processes.

Interpreting infant behaviour, including eye movements, can be compared with interpreting ancient cave paintings.

This makes it essential to come up with good explanatory theories able to be tested by **implementing** them.

Other tests are possible, though nothing is decisive, e.g.:

- Seeing how they react to unusual obstacles impeding their goals.
- Seeing how they react to unusual tools, opportunities, toys.
- Seeing how they react to adults who (deliberately) make mistakes in actions or in answers to questions.
- Seeing how they help someone trying and failing to do something.
- Asking them what would happen if... in various situations.
- Seeing how they react to fallacious arguments e.g. a toy wooden object will grow if planted in the garden because wooden things grow in the garden.
- Asking them to do something like pick up a shadow, get into a small toy car, wash the floor with a thimbleful of water and a toothpick,


Some of these methods can also be used with other animals, and future robots

Testing understanding of gears

Kinds of Causation: (Humean)

Two gear wheels attached to a box with hidden contents.

Can you tell by looking what will happen to one wheel if you rotate the other about its central axis?



- You can tell by experimenting: you may or may not discover a correlation – depending on what is inside the box.
- In more complex cases there might be a knob or lever on the box, and you might discover a dependence: the position of the knob or lever determines how the first wheel's rotation affects the second wheel's rotation.
(Compare learning about gears in driving a car.)
- In still more complex cases there may be various knobs and levers, modifying one another's effects through hidden mechanisms. There could also be motors turning things in different directions, competing through friction devices, so that the fastest one wins.

Meshed gear wheels are different

Kinds of Causation: 2 (Kantian)

Two more gear wheels:

You (and some children) can tell, by looking, how rotation of one wheel will affect the other.

How? You can simulate rotations and observe the consequences.

(Making assumptions about the kind of stuff the wheels are made of: [rigid](#) and [impenetrable](#))

What you can see includes this:

As a tooth near the centre of the picture moves up or down it will come into contact with a tooth from the other wheel.


If the first tooth continues moving, it will push the other in the same direction, causing its wheel to rotate.

(I am not claiming that children need to reason verbally like this, consciously or unconsciously.)

NB: The simulations that you run can make use of not just perceived shape, but also [unperceived constraints](#): in this case rigidity and impenetrability.

These need to be part of the perceiver's ontology and integrated into the simulations, for the simulation to be deterministic.

The constraints and processes using them need not be conscious, or expressed in linguistic or logical form: how all this works remains to be explained.




Reasoning about partially observable mechanisms

Kantian causal explanations

Reasoning that is used to predict what will be observed can also be used to explain what is observed, if mechanisms are partly hidden.

A similar kind of reasoning can be used to explain observed motion: e.g. if the left wheel is moved as shown by the left arrow, why does the right wheel move as shown by the right arrow?

There is no way to infer the correct explanation by reasoning from what is observed,
But if an appropriate **theory** can be constructed, and added as a “hidden premiss” (abductive reasoning) then it can be tested by working out its consequences and comparing them with what is observed.



Humean (nowadays Bayesian) causal reasoning based on statistical information about perceived correlations, is the only kind of causal reasoning available **when mechanisms are not understood**.


Kantian (structure-based, deterministic) causal reasoning is sometimes available, and is often more useful, but it requires a richer ontology and more general reasoning abilities that can make direct use of structural constraints in complex configurations.

Testing a child's ability to do such reasoning

Perhaps the following configuration could be used to test whether a child had moved beyond the empirical understanding of gear-wheel interaction.

**If wheel A moves as shown,
which way will wheels C and D turn?**

This can test adults too!



Kant implicitly supports something like our theory in his question

For how should **our faculty of knowledge** be awakened into action did not objects affecting our senses ... work up the raw material of the sensible impressions into that knowledge of objects which is entitled experience?

Whatever this “faculty of knowledge” is, it could not be **awakened** without **already existing**.

What is this faculty?

Perhaps it is a kind of information-processing architecture that drives exploration of the environment while it grows itself?

Compare the following sketchy account of cognitive and meta-cognitive epigenesis:

Some virtual machines can extend themselves indefinitely

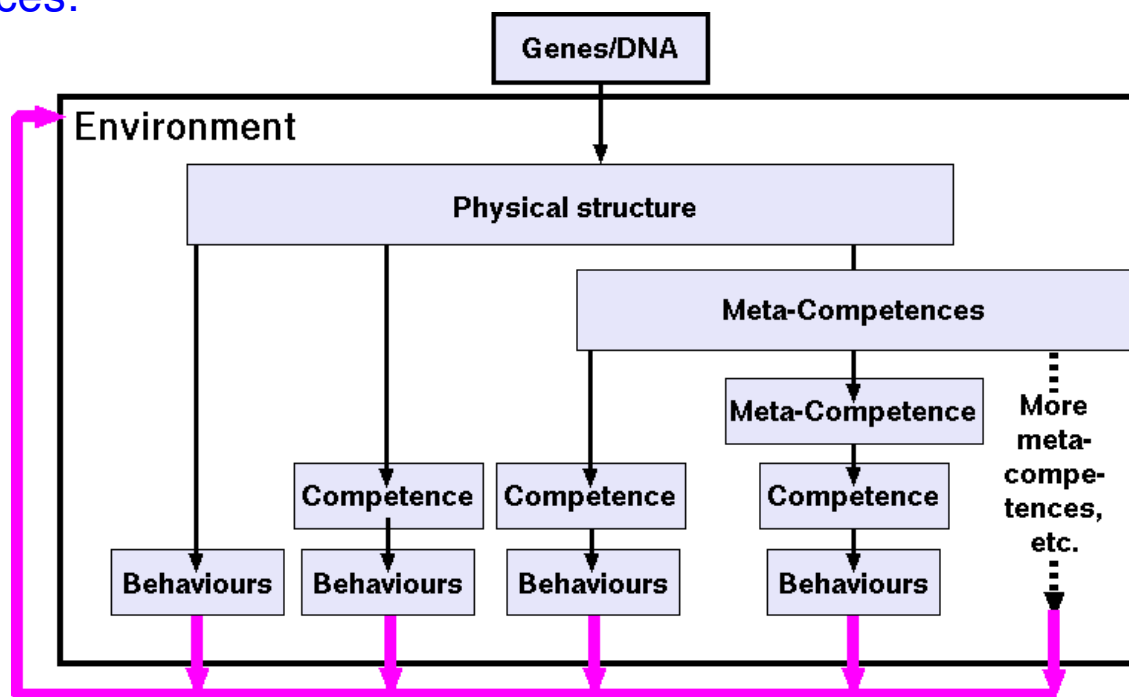
Cognitive epigenesis: Multiple routes from DNA to behaviour, some via the environment

Meta-configured meta-competences:

Humans not only learn new kinds of things, they can learn to learn even more varied kinds of things. E.g. after completing a degree in physics you are enabled to learn things that a non-physicist could not learn, e.g. learning how to do more sophisticated experiments.

Meta-configured meta-competences:

(towards the right of the diagram) are produced through **interaction** of pre-configured or previously produced meta-configured competences with the environment, including possibly the social environment



<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>

Natural and artificial meta-configured altricial information-processing systems
Chappell&Sloman, 2007, IJUC

Diagrams developed with Jackie Chappell and Chris Miall.

Infinite cognition (competence, not performance)

How can a finite child come to think about infinite mathematical objects?

It was suggested above that learning about numbers involved learning to set up and make use of one-to-one mappings between structures and/or processes (e.g. counting processes) – especially synchronising two concurrent discrete processes.

The processes can have different stopping conditions, depending on the task:

Different stopping conditions are required for the two processes:

- (a) measuring a distance of 25 paces by pacing, and counting;
- (b) measuring the diameter of a field by pacing and counting;

Stop (a) when the counting reaches 25, and stop (b) when the pacing reaches the far edge of the field.

In case (a), stopping is controlled by the counter, and the final number known in advance, whereas in (b) the stopping is controlled by the world.

Since stopping in (b) is up to the world, the world might never satisfy the stopping condition.

Understanding the questions: “Will the process stop?” “When will the process stop?” is part of understanding that a process has a stopping condition and that it needs to be monitored to determine whether the condition has been reached.

Recognising that some conditions may never be reached (e.g. stop when you return to your starting point) can be the basis of understanding that a sequence can be infinite.

Lance Rips and colleagues (BBS 31,6 2008) seem to suggest that the genome provides something like the axiom of induction.

I suggest that the need to accommodate infinite numbers can be inferred from the fact that some counting sequences may not stop.

Could there be a largest number empirically

It is an empirical fact that I have only ten fingers: the world stops my counting when I count my fingers (properly).

If the universe had had only a finite number of countable things would that have made the set of positive integers finite?

No because counting is not restricted to actual objects: we can count possible things too.

Initially a child learns a fixed set of arbitrary number names (numerals), e.g. “one” to “six”, then goes on extending it, e.g. up to “twenty”.

Later, as a result of our cultural evolution, the child learns one or more generative schemes for producing a determinate but never-ending sequence of number names.

This sequence can be used not only for counting physical objects or events, but also for counting subsets of the number names, e.g. the numbers between 23 and 37.

So if following the instruction to keep generating the next number name and to count all the number names, the child could never stop.

A reflective child could notice this feature of the process of counting the counting process.

This does not depend on the use of any particular generative scheme, any scheme that provides a rule for generating the next number after any given number will do.

We can tell that we have such a scheme, and therefore that there cannot be a largest number: which provides a sufficient basis for guaranteeing the existence of at least the simplest kind of infinity.

The number zero arises from cases where you cannot even start counting! (Noticing it took centuries.)

WARNING

This presentation should not be treated as an authoritative account of the ideas of any of the thinkers mentioned herein.

My summaries of the views of others are often interpretations of things I have read recently or things I remember reading in the past. I am interested in them as theories that could be held by serious thinkers, and don't really care who did or did not hold them.

Anyone seriously interested in the views of Kant, Hume, Frege, Russell, Feynman, etc. should study their writings instead of relying on my summaries.

A useful overview of ideas about the nature of mathematical discovery and mathematical reasoning around the time of Kant can be found in

James Franklin, *Artifice and the Natural World: Mathematics, Logic, Technology*, in *Cambridge History of Eighteenth Century Philosophy*, Ed. Knud Haakonssen
Cambridge University Press, 2 volumes 1423 pp. 2006

Franklin's chapter is available online without footnotes here:

<http://web.maths.unsw.edu.au/~jim/18c.html>

It is interesting that he too notes the possibility of testing ideas in philosophy of mathematics by implementing working intelligent systems:

Hume's views on inference are seen to better advantage if they are thought of not in terms of formal logic, or even introspection, but as a research proposal to be implemented in, say, silicon chips. Modern Artificial Intelligence, like most eighteenth century writing, is concerned with the implementation of a system of inference, not just the formal structure of the system itself. From that point of view, it is necessary to answer questions that do not arise in formal logic, such as how the symbols become attached to the things they mean.

On the last point see also this critique of symbol grounding theory:

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

See his survey of philosophies of mathematics: <http://web.maths.unsw.edu.au/~jim/philmathschools.html>

(X) Related presentations

This is one of a series of (overlapping) presentations linking biological evolution, developmental psychology, robotics and mathematics.

The first two, given in 2008 were:

Could a Child Robot Grow Up To be A Mathematician And Philosopher?

Talk at University of Liverpool, 21st Jan 2008

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks#math-robot>

Includes several examples of problems with both empirical and mathematical aspects.

Could a Baby Robot Grow up to be a Mathematician and Philosopher?

Shorter version, presented at 7th International Conference on Mathematical Knowledge Management (MKM'08). University of Birmingham, 29 Jul 2008

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#mkm08>

Later versions (available at the same web site) are aimed more at psychologists and biologists interested in human cognitive development and the mechanisms that make it possible.

I attempt to present some important relationships between interaction with the environment and the development of mathematical competence, but in a simplified way, starting with an example that is intended to give adult researchers a feel for some of the problems confronting a young child learning about a new range of phenomena, and gaining new insights into previously acquired knowledge.

The ideas here are developments of work in my 1962 DPhil thesis (now online), a paper on varieties of knowledge in 1965, a paper on analogical reasoning in IJCAI 1971, Chapter 8 of The Computer Revolution in Philosophy (1978). It builds on and overlaps with the two presentations above, but has a new introduction, and attempts to give a deeper analysis of the kinds of changes that can occur during learning.

Some are related to nature-nurture issues discussed in Chappell & Sloman (2007) "Natural and artificial meta-configured altricial information-processing systems" <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>

Related online papers and presentations

Functions of vision, with speculations about the role of multiple multistable dynamical systems:

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#compmod07>

Discussion paper on Predicting Affordance Changes (including epistemic affordances)

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0702>

For a critique of the notion that reduction to logic can explain mathematical knowledge (similar to Frege's argument) see my DPhil thesis (Oxford, 1962), e.g. Appendix II

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#706>

Knowing and Understanding: Relations between meaning and truth,
meaning and necessary truth, meaning and synthetic necessary truth

The arguments were partly replicated here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#712>

Explaining Logical Necessity *Proc. Aristotelian Society*, 1968/9, Vol, 69, pp 33–50.

Then linked to AI theories of representation in 1971

<http://www.cs.bham.ac.uk/research/projects/cogaff/04.html#200407>

Computational Cognitive Epigenetics (With J. Chappell in BBS 2007):

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703>

Commentary on Jablonka and Lamb (2005)

See other papers and presentations with J. Chappell here:

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/>

Our presentations on causation at WONAC, Oxford June 2007:

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac>

An argument that internal generalised languages (GLs) preceded use of external languages for communication, both in evolution and in development:

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

What evolved first: Languages for communicating, or languages for thinking
(Generalised Languages: GLs) ?

More

Diversity of Developmental Trajectories in Natural and Artificial Intelligence (AAAI Fall symposium 2007):

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0704>

Actual Possibilities (in KR 1996) <http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15>

Virtual Machines in Philosophy, Engineering & Biology (at WPE 2008):

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#wpe08>

Ideas about how machines or animals can use symbols to refer to unobservable entities
(why symbol grounding theory is wrong)

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

Introduction to key ideas of semantic models, implicit definitions and symbol tethering

Additional papers and presentations:

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/>

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/>

<http://www.cs.bham.ac.uk/research/projects/cogaff/>

Shaaron Ainsworth drew my attention to the work of David Tall after many of these ideas had been developed.

See <http://www.warwick.ac.uk/staff/David.Tall/>

There seems to be a lot of overlap, although our main focus is different, insofar as I am trying to find explanations of how the cognitive mechanisms work, how they evolved, how they develop etc. I think his concept of “procept” is closely related to the notion of procedure, algorithm or program, as used in Computer science and AI as something that can both be used (i.e. run) and referred to or manipulated. Not all programming languages support the dual role, as AI languages began to do in the 1960s, e.g. Lisp, and later Pop2, Pop-11, Prolog, Scheme, ... also Algol68. All these built on Church’s Lambda Calculus, which built on Frege’s ground-breaking work on functions. It would be interesting to investigate how Tall’s independently developed ideas are similar and different. G.J. Sussman (see HACKER) and others showed how such notions are crucial for certain kinds of learning. See also Chapter 8 of Sloman 1978, on learning about numbers.

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

NOTE: Although I don’t give any references to G.Polya, the approach taken here is much influenced by reading his *How to Solve It*, many years ago.

Other useful references

James Franklin, Adrian Heathcote, Anne Newstead: 'MANIFESTO: Mathematics, the science of real structure: An Aristotelian realist philosophy of mathematics,'

<http://web.maths.unsw.edu.au/~jim/manifesto.html>

James Franklin, Last bastion of reason *New Criterion* 18 (9) (May 2000), 74-8. (Suggests Lakatos is dishonest about mathematics.)

<http://newcriterion.com:81/archive/18/may00/lakatos.htm>

Teri Merrick, 'What Frege Meant When He Said: Kant is Right about Geometry', *Philosophia Mathematica* Vol 14, 1, 2006. doi:10.1093/philmat/nkj013

Lisa A. Shabel, "Kant's Philosophy of Mathematics" in *The Cambridge Companion to Kant*, CUP, 2006.

<http://people.cohums.ohio-state.edu/shabel1/cv.html>

Jeanette Wing on 'Computational thinking' *CACM* 49,3, March 2006, pp. 33-35.

<http://www.cs.cmu.edu/afs/cs/usr/wing/www/publications/Wing06.pdf>

Anette Karmiloff-Smith *Beyond Modularity*. See the BBS precis:

<http://www.bbsonline.org/documents/a/00/00/05/33/index.html>

Douglas Hofstadter (Thanks to Dan Ghica for the link.)

<http://www.stanford.edu/group/SHR/4-2/text/hofstadter.html>

"on seeing A's and seeing As", *Stanford Humanities Review* Volume 4, issue 2

Mateja Jamnik *Mathematical Reasoning with Diagrams: From Intuition to Automation* (CSLI Press 2001)

<http://www.cl.cam.ac.uk/~mj201/research/book/index.html>

Daniel Winterstein *Using Diagrammatic Reasoning for Theorem Proving in a Continuous Domain*

PhD thesis, Edinburgh, 2004, <http://winterstein.me.uk/academic/>

Alison Pease, *A Computational Model of Lakatos-style Reasoning* PhD thesis, Edinburgh, 2007

<http://hdl.handle.net/1842/2113>

Samson Abramsky & Bob Coecke, *Physics from Computer Science* IJUC, vol. 3:3, 179–197, 2007 (QM in diagrams)

<http://web.comlab.ox.ac.uk/oucl/work/samson.abramsky/YORKIJUC.pdf>