Talk presented at Dagstuhl Workshop 28th March 2011
Modified Version at Oxford Workshop 6th April 2011
http://www.comlab.ox.ac.uk/automatheo2011/program.html

Evolution, robots and mathematics

or

Some largely unnoticed requirements for future intelligent robots

(Partly based on work by Piaget, Karmiloff-Smith and others)

Aaron Sloman
http://www.cs.bham.ac.uk/~axs/

I suggest some goals for research on intelligent robots based partly on a "rational reconstruction" of Piaget's theories, and Annette Karmiloff-Smith's ideas in Beyond Modularity (1992). (AK-S)
I suggest this could transform both developmental psychology and AI.
These slides (liable to ongoing development) and related slides are available in my 'talks' directory:

http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk90

Annotated videos used are here:

http://www.cs.bham.ac.uk/research/projects/cogaff/movies/vid

A related presentation will be given at AGI 2011: http://agi-conf.org/2011/

Philosophy of mathematics, robotics and biology

There are philosophical theories about the content of mathematical knowledge and its status – in both cases compared with other kinds of knowledge, e.g. how mathematical discovery and knowledge differs from knowledge about physics, geology, history, mental contents, ...

I originally got into AI (around 1971) mainly through an interest in philosophy of mathematics, and in particular an attempt to defend Kant's view of mathematics as both non-empirical (apriori, but not innate, since discoveries are made) and also synthetic (non-trivial, and also not purely logical). (Compare: Plato, Hume, Russell, Hilbert, and others.)

The aim was to show why Kant was (roughly) right – by showing how a suitably designed baby robot could grow up to be a mathematician, by making its own mathematical discoveries, e.g. about geometry, arithmetic, topology, sets, etc.

The mechanisms could also shed new light on developmental psychology and attempts to understand what biological evolution had achieved. E.g. What's in human genomes?

Work reported here is a continuation of that investigation: **a long way to go still.**
For more on the philosophy see related presentations on the same web site, e.g.

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks#math-robot`
Could a Child Robot Grow Up To be A Mathematician And Philosopher?

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddlers`
A New Approach to Philosophy of Mathematics: Design a young explorer, able to discover "toddler theorems"

Key idea: "toddler theorem"

Infants and toddlers discover much through play and exploration –
then discover (unconsciously) that they can work some things out instead.

Examples:

- How to avoid pain when using a hand to shut a drawer or door
- Pulling an inelastic straight string makes the far end move
- Pushing a long rigid object makes the far end move
- Rotating a rigid object held at one end makes the far end move (and move more)
- Moving a spoon containing stuff horizontally makes the stuff move horizontally
- Containment is transitive. (Useful when feeding yourself.)

  OLDER CHILDREN

- Counting fingers left to right and right to left gives same result
- 10 or 12 blocks can be arranged in a rectangle, but not 11 or 13 blocks (only in a line) (see later).

There are further stages of knowledge extension and meta-cognition, not discussed here.
Some aspects are discussed in Karmiloff-Smith (1992).

Adults can also play and learn like toddlers

Exploring new domains.

Various board games (each defines a domain of possible positions and legal moves)

What can you do with a rubber band and pins?

How can you arrange a collection of blocks of different heights in descending order without having to back-track?

It's one thing to find and be able to use an efficient strategy.
It's another thing to be able to describe the strategy (e.g. so that others can use it).
Further competences are required for proving that the strategy works.
Yet more sophistication is required to compare alternative strategies and investigate whether one is optimal.
Sometimes possibility for generalisation or parametrisation is noticed.

Similar comments about stages of understanding can be made about many more examples in this presentation.

AK-S noticed some of this, but seems to be unaware of details that a mathematician would notice.

Which rows can be rearranged to form rectangles?

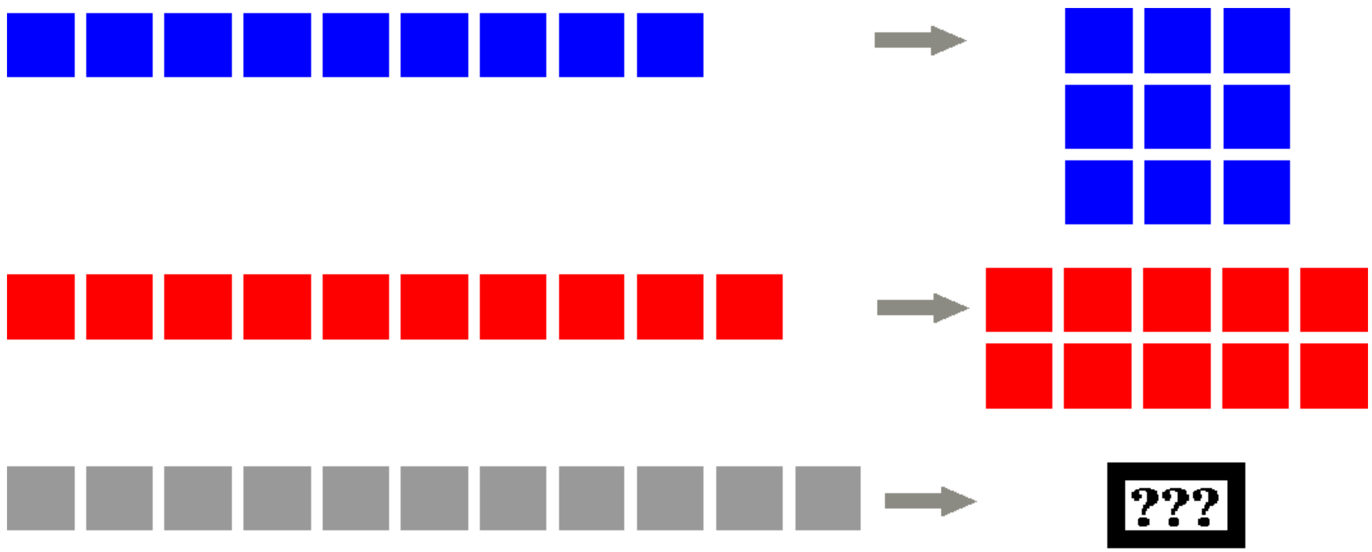

Why can some rows of blocks be rearranged as rows and columns but not all?

Does it depend on the colour of the blocks? Or what they are made of?

Why not?

How can a learner be sure the third row cannot be rearranged into a regular array of rows and columns?

Not all animals, and not all humans, learn the same things.

(Age vs genetic vs environmental/cultural reasons.)

Could this be extended to a geometric proof of the unique factorisation theorem?

Doing science: Don't just study one example

Trying to understand cognition, or intelligence by studying one species is like trying to build a science of chemistry on the study of just one very complex molecule.

To understand a complex functioning system you need to know what the consequences would have been if various features had been absent, or if other features had been present.

That sort of science needs to be inherently comparative.

Evolution has provided many different samples from the space of possible designs.

- I find it useful to spend a lot of time talking not only to developmental psychologists studying humans but also biologists studying non-human animals.
- In order to understand the nature of mathematical thinking, discovery and proof, it's not enough to study only what adult mathematicians think, say and do.

  Insights can be gained from some of the intermediate stages in development of future mathematicians: e.g. infants and toddlers.
- Unfortunately, researchers don't always know what to look for, and may not notice important details, especially if
  - they don't have deep personal experience of doing mathematics
  - they are unfamiliar with the diversity of philosophical theories about the nature of mathematics.

What's common to the portia spider crows building nests orangutans moving through tree canopies and humans building mathematical theories?

What's common to the portia spider crows building nests orangutans moving through tree canopies and humans building mathematical theories?

The evolutionary impact of an environment containing 3-D structures, processes and kinds of stuff.

It's a long, complicated and still largely mysterious story...

It's mostly a story about requirements or problems faced by organisms.

A related story concerns the designs that meet those requirements

designs produced by evolution, often in combination wth learning and the learner's environment: designs that are solutions to different subsets of the requirements/problems.

Regarding Portia spiders, see (Tarsitano, 2006)

`http://dx.doi.org/10.1016/j.anbehav.2006.05.007`

From the abstract:

"The jumping spider Portia labiata can complete detours in which it must move away from a goal (i.e. prey) before approaching it. This detouring behaviour can be divided into two phases: a scanning phase, during which Portia stays roughly in one spot and examines its environment using its principal eyes, and a locomotory phase, during which Portia performs the detour. ..."

John McCarthy in "The Well-Designed Child" (1)

Some researchers assume humans are born with no knowledge, only a powerful learning engine and a large empty memory
– and a robot could do likewise. Compare (Turing, 1950).

McCarthy thought otherwise (McCarthy, 2008).

"Evolution solved a different problem than that of starting a baby with no a priori assumptions."

....

"Instead of building babies as Cartesian philosophers taking nothing but their sensations for granted, evolution produced babies with innate prejudices that correspond to facts about the world and babies' positions in it.

Learning starts from these prejudices.

What is the world like, and what are these instinctive prejudices?"

John McCarthy in "The Well-Designed Child" (2)

"Instead of building babies as Cartesian philosophers taking nothing but their sensations for granted, evolution produced babies with innate prejudices that correspond to facts about the world and babies' positions in it.

Learning starts from these prejudices.

What is the world like,
and what are these instinctive prejudices?"

Different answers

– for different species

– for different stages of development within a species

foetus, neonate, infant, toddler, child, youth ...quantum physicist...

“What is the world like?”
Not for a physicist, but for a baby animal or robot
(depends on the species)

It’s a tangled web of partial orderings, possibilities and necessities:

of size, volume, distance, angles, directions, speeds, forces, pressures, dynamical properties, temperatures, kinds of stuff, kinds of motion, kinds of causation ...

Embedded in a unifying topology

containment is transitive

changes can be continuous

with no global metrics (no cartesian coordinates)

but semi-metrical extensions of partial orderings

E.g. the distance between A and B, exceeds the distance between C and D, by less than the distance between E and F.

Different subsets can be sensed by different species
or by the same species at different stages of development.

NB. This is not “a blooming buzzing confusion” as suggested by William James.

See this draft discussion note
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/blooming-buzzing.html`

Separating out domains

Although the spatio-temporal environment can be thought of as a huge, unstructured, mess (as suggested by James) there is actually a great deal of structure that has been found over centuries by craft-workers, engineers and scientists.

Some of that structure is represented by different scientific/engineering disciplines.

Conjecture:

For a human infant or toddler some of the structure can be found by dividing the environment into different domains, distinguished in part by performing different sorts of actions with systematic variations, and discovering their effects.

An infant, or someone trying to help the infant develop, can identify different domains of action and carve out subsets of space-time.

Key themes

- Trying to understand cognition by studying ONE species, or building ONE robot is like trying to develop a science of chemistry by studying ONE complex molecule.

  Science needs to be **comparative** in order to achieve generality and depth.

- Understanding complex functioning systems is impossible without understanding the requirements that their designs satisfy (or partly satisfy).

  That is also a comparative study: biological designs can (partly or wholly) satisfy different requirements, and the same requirements can be (partly or wholly) satisfied by different designs.

- Instead of only requirements and designs, some types of organism have meta-requirements and meta-designs where those requirements and designs can be instantiated differently in different contexts.

  Compare parametric polymorphism in OOP ?

- In biology the instantiation is not just initial instantiation: it can be staggered over time – some parameters being acquired later than others (e.g. taste in food, social norms)

  In some altricial species the parameters are not just provided by the environment, but actively sought and constructed by individuals

- In humans (and future robots) that instantiation process is not monotonic/ incremental learning – it includes growth and reorganisation of major features.

  In humans language learning is a well known example, including “U-shaped” learning. I suggest this is a special case of something deeper and more general.
  Compare Karmiloff-Smith (1992) (Beyond Modularity).

What should a bright young animal learn? (1)

Do not assume that every learner starts off

– provided with vectors of scalar values from sensors

and

– ability to generate vectors of scalar values to be sent to effectors

– with the problem of learning statistical patterns relating them.

Would something like that enable a foal to stand up, find its mother's nipple and suck, and then run with the herd within hours or minutes?

There's plenty of evidence that biological evolution can provide powerful usable competences (presumably mainly encoded in the genome and its products) available at birth.

When there don't seem to be such competences perhaps that's because we don't know what to look for.

E.g. one option is a large collection of highly specific **parametrised** competences waiting for their parameters to be discovered.

(Some versions of Chomsky's nativism.)

More complex possibilities are discussed in (Chappell & Sloman, 2007)

What should a bright young animal learn? (2)

Our universe is very rich, with discoverable structure of various kinds on many spatial and temporal scales.

The information available to a young learner is not a formless mush, but can be actively partitioned, into subsets (“domains”, “microworlds”).

What should a bright young animal learn? (2)

Our universe is very rich, with discoverable structure of various kinds on many spatial and temporal scales.

The information available to a young learner is not a formless mush, but can be actively partitioned, into subsets (“domains”, “microworlds”).

For a toddler, there are small spatial objects that can be pushed, lifted, arranged in various ways on various surfaces, stacked, thrown, prodded, banged together, banged against other things, hit with other things, rubbed together, rubbed on other things, sucked, swallowed, ...

(Show videos: blanket, piano, yogurt)

Given a collection of such objects there are various ways in which exploratory actions involving them can be performed,

e.g. doing various things with one of them; doing things by combining them, adding one at a time; rearranging configurations by moving items one at a time, ..., and more.

An animal with a sound generator and auditory receptor can explore different sound patterns – some produced using things in the environment.

If it can measure time it can try doing things at different speeds, or with varying speeds.

There are many more examples.

For human infants, see for example (Rochat, 2001),
and Spelke’s work cited in several contexts in Karmiloff-Smith (1992).

Analysing types of exploration domain

Much research by developmental psychologists has focused on specific competences – perceptual, motor, attentional, etc.

As far as I know there has been no attempt to produce a systematic survey of the various domains that can be explored by young animals including

- analysing their links with sensory and motor capabilities
- characterising their mathematical structure

  (including discrete, continuous, quantitative, relational, parametrisation, ...)
- dissecting possible exploration trajectories they support
- studying the ways they can be combined with other domains (e.g. through spatial juxtaposition, temporal/sequential combinations, etc.)
- looking for generative specifications

  E.g. axiomatisations, expression as a grammar for possibilities, ad-hoc assemblages and composition rules, (precursors to grammars).

Types of domain (extending Karmiloff-Smith (1992))

- Some exploration domains are continuous, some discrete.
- Some essentially involve a particular type of material (skin, paper, wood, metal, water, plasticine, wire, etc.)
- Some involve spatial features – shapes, relations, pathways, and/or temporal patterns, independently of any particular kind of material (e.g. sum of angles of a triangle)
- Some are deterministic, e.g. motions of tokens on a 2-D grid, effects of sorting operations on items that can be ordered, repeating patterns.
- Others are (at least from the point of view of the explorer) non-deterministic

  e.g. because of unpredictable variations in continuous motions or because discrete events are generated by inaccessible or unknown mechanisms. (Dice, bagatelles.)
- Two special kinds of knowledge about a domain (see Piaget's last two books):
  - Knowledge of possibilities

    What configurations of objects or events or processes **can** occur
  - Knowledge of constraints/necessities

    Which relationships, arrangements **are impossible**,
    and what the **necessary consequences** of certain arrangements are:
    e.g. if A contains B and B contains C then A must contain C.
  - Knowledge of probabilities

    (I think the role of this has been over-inflated by learning theorists)
- Some domains include intelligent agents, with percepts, desires, intentions, etc.
  These require architectures able to deal with referential opacity. (Meta-semantic competences.)

Carving the world up into exploration domains

Biological evolution gives some species (e.g. deer that need to run with the herd at birth) powerful fully formed competences – suited to a particular range of environments. (Precocial/preconfigured competences)

For other species most competences are not precocial or preconfigured in the genome, but meta-configured: the genome provides means of acquiring the the competences needed in different environments

See (Chappell & Sloman, 2007).

In many cases (e.g. in young humans) this requires both physical capabilities and perceptual and control mechanisms that will support separating out realms of possibility (exploration domains).

This can be done by systematically trying different sorts of experiments, on the same objects;
and also trying the same experiments on different sets of objects.

In simple cases such experiments include grasping, releasing, pulling, pushing, sucking, chewing, twisting, throwing, banging, and doing various things to other objects with one that is held.

For examples see these annotated videos of children playing especially the piano and yogurt videos:
`http://www.cs.bham.ac.uk/research/projects/cogaff/movies/vid`

Piaget's two last books

These two books by Piaget and colleagues do not seem to be widely known:

Piaget, Jean, et al.,
Possibility and Necessity Vol 1. The role of possibility in cognitive development,
U. of Minnesota Press, Tr. by Helga Feider from French in 1987, (Original 1981)

Piaget, Jean, et al.,
Possibility and Necessity Vol 2. The role of necessity in cognitive development,
U. of Minnesota Press, Tr. by Helga Feider from French in 1987, (Original 1983)

As far as I know there has been no development of the ideas in these two books since their publication.

Before I learnt about these books, I consulted several well known researchers in developmental psychology about some of the research problems mentioned below. None mentioned Piaget's books – until I asked Annette Karmiloff-Smith, to whom I am very grateful for the pointer.

Reading these books may be hard for people unfamiliar with Piaget's work or developmental psychology – though mathematicians and physicists may find it easier than others.

An older book which I found very inspiring many years ago was partly stimulated by Piaget's work:
The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology.
Sauvy and Sauvy (1974)

Annette Karmiloff-Smith's Beyond modularity (1992), presents additional material and in some ways takes the ideas further than Piaget did. Her book should be compulsory reading for cognitive roboticists.

See also Borovik (2010).

I may later add other links to relevant online material on Piaget. Suggestions welcome

Karmiloff-Smith (1992) on domains

AK-S makes important claims that I recently discovered overlap with what I have been claiming about biological precursors to mathematical competences. (Sloman, 2010) See http://www.cs.bham.ac.uk/research/projects/cogaff/talks/

- Children (and other animals) can explore a domain and acquire both factual knowledge about it and behavioural mastery

  E.g. learning to communicate using words, phrases, sentences and ways of putting them together.

- Later, without being aware of doing so, they reorganise that knowledge into a more systematic form, e.g. starting to use more powerful/general syntactic rules.

  The new format can be referred to as a deductive theory, since it allows new possibilities and new necessities to be discovered by reasoning, i.e. non-empirically. (It need not use logic.)

  Learners cannot express this new knowledge verbally, but show it in their behaviour

  e.g. what they treat as possible or impossible.

- Later still, in some cases, they become able to express this new knowledge explicitly (e.g. answering questions) though they may get some things wrong.
- Later still, in some cases they acquire explicit mastery of the domain, being able to reason correctly about it.

  Although she does not say this, if the learner also has the meta-semantic competence required to represent and reason about what other agents do and do not know that can make possible explicit teaching, advice, assistance, etc.

- AK-S describes these processes as "representational redescription",

  but there must be more than change of representation: new ontologies, new algorithms and new architectures may be required.

A richer variety of transitions: Not yet robots

- Purely exploratory actions and observations

  Need to specify much more detail here
- Examples stored for re-use
- Examples parametrised to form patterns with different instances
- More complex structures and processes: e.g. reach then grasp then lift then rotate then ungrasp
- Differences noticed (ungrasp over table, ungrasp over floor)
- New concepts created for kinds of stuff, kinds of relation, kinds of process, kinds of combinations of structure and process (e.g. ”stopping motion” ”preventing motion”).
- **Abductive transformation: relatively small number of re-usable general principles found**
  A framework theory constructed, unifying a domain, and allowing prediction in novel circumstances.
- Metacognitive mechanisms grow and gain access to the new theory

  Child can think up new possibilities, notice that there are constraints, answer questions – some incorrectly.

  The meta-theory is gradually extended and debugged
- If child treats others as agents can start sharing this expertise – collaborating, demonstrating, warning, admonishing, correcting, ...
- Later the child learns a human language and the formalizations become richer and more general (e.g. using quantifiers, modal operators, etc.)
- The linguistically formulated meta-theory may be inaccurate and will need to be debugged: it too could start off as a patchwork then later become reorganised into a framework theory.
- Later domains are combined and possibly unified if they share some structure.
- **Different learners will inevitably have different learning trajectories.**

Conjecture: Before human language...

Much knowledge (not all) starts Empirical then becomes proto-mathematical

Language learning: the “U”-shaped curve

- Language learning initially uses stored, re-usable, learnt patterns:
- Then becomes rule-based
  - When this happens the rules are at first over-generalised causing production errors that were not present previously.
  - This is a temporary phase.
- Then language use becomes rule+pattern based
  - As a result exceptions to the syntactic rules are handled properly.
  - This requires a significant change to the information processing architecture (which would not be easy to program).
- In parallel with that development, mechanisms develop that allow many stored patterns to be retained for rapid comprehension and production (non-trivial extension)
  - This can dominate processing in experts most of the time.
  - (This can also lead language researchers to draw false conclusions about the nature of linguistic competences – e.g. assuming they are inherently statistical.)

The Conjecture (based on work with Jackie Chappell)

The case of language is a later evolutionary development of something much older that has largely gone unnoticed in other animals and pre-verbal children (Sloman, 2008a)

The conjectured genetically determined mechanisms support “layered learning”, where what can be learnt in new layers depends on what was previously learnt.

The genetically pre-configured mechanisms evolved in response to very general features of our world, common to explorable domains in many different environments

e.g. different domains including 3-D spatio-temporal structures and processes including topological and geometrical states and processes involving different kinds of stuff (matter) with different (unobservable) capabilities (fragility, elasticity, thermal conductivity, toxicity, etc.)

But what structures the mechanisms actually allow to grow in the individual’s information processing architecture can depend on

The morphology and sensorimotor and other mechanisms in the species.

The kind of environment in which exploring and acting occur

Moreover, I conjecture that, like language learning (which I think evolved later), these more general mechanisms start only pattern-based, then after a while (after enough pattern information has been acquired empirically to be worth compressing) the new meta-cognitive mechanisms attempt to construct a generative/productive theory that can cope with novel situations and generate new patterns.

They don’t use a totally general highly optimised set of mechanisms, but satisficing mechanisms (or meta-mechanisms) selected by evolution – to fit a range of space-time environments (McCarthy, 2008).

A diagrammatic summary

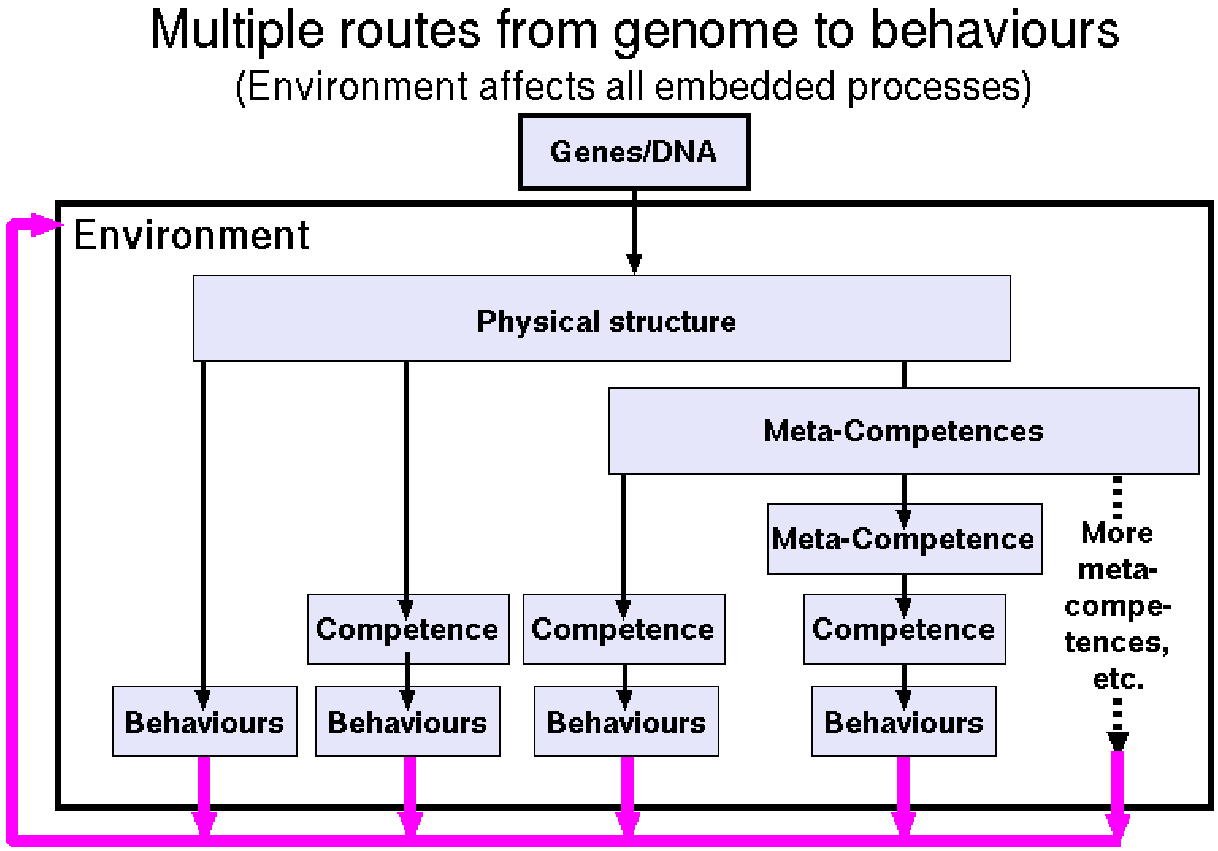

Work done in collaboration with Jackie Chappell (Chappell & Sloman, 2007)
Chris Miall helped with the diagram.

These ideas are still too vague to be implemented: the point added in this presentation is that many of the competences and metacompetences are concerned with the sorts of exploration domains described above, and ways of (a) combining them into more complex domains (b) finding deeper, hidden, explanatory domains.

Different “domains” of development

Different kinds of knowledge, with associated competences

Karmiloff-Smith’s microdomains

AI microworlds

Things to learn about

e.g.

- kinds of action – different modalities.
- kinds of perception (different modalities)
- Kinds of environmental structure/process/material

  E.g. for humans
  - kinds of food
  - kinds of clothing
  - kinds of physical interaction with adults
  - kinds of toys
  - kinds of stuff (material)
  - kinds of structure or process
  - social interactions
  - language

NB: space-time allows different domains to be combined.

including structured actions, structured objects

Transitions in information-processing requirements

Changing relationships between organism and environment

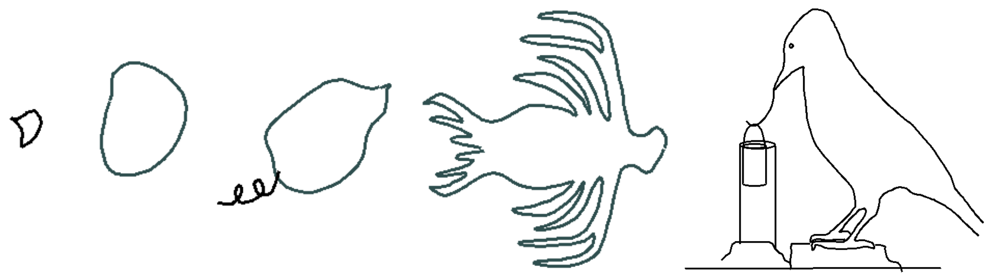

Types of environment with different information-processing requirements

- Microbes in chemical soup, can be wholly dependent on nutrients in their neighbourhood.
- Soup with detectable gradients: offers opportunities to improve location if motion can be controlled.
- Soup plus some stable structures, e.g. places with good stuff, bad stuff, obstacles, supports, shelters, offers advantages for organisms that can build up long term spatial memories, and can plan routes.
- Things that have to be manipulated to be eaten (e.g. disassembled) need new forms of process perception and process control, including ontologies that include kinds of "stuff" with different properties.
- Controllable manipulators: mouth, hands, feet, require different uses of information.
- Things that try to eat you – lead to forms of control for escaping, hiding, defending, etc.
- Food that tries to escape – lead to forms of control for chasing, trapping, lying in wait, heading off, etc.
- Mates with preferences, lead to forms of control of behaviours for attracting attention and winning favour.
- Competitors for food and mates lead to deception, fighting, warding off, defending territory, etc.
- Collaboration with others requires action controlled to aid collaboration, and possibly communication.
- and so on (including information-processing in plants) .....

How information-processing **requirements** change across such cases, depends on both features of the environment and features of the organism (products of previous evolution). Contrast: how **designs** change.

The methodological significance of transitions

The notion of “transition”, applicable in several different contexts and at different levels of abstraction, can help us understand the varying relationships between genomes and information-processing architectures.

Maynard Smith and Szathmáry discussed eight types of transitions in evolution. But there are many more relevant to the evolution of information processing capabilities (minds):

- transitions in the niches, or sets of requirements that play a role in natural selection, via changes in the environments in which organisms live, compete and die
  (including changing information available/unavailable about constraints and opportunities);
- transitions in the information-processing competences and designs that are relevant to coping better or worse with those requirements;
- transitions in the implementations of those designs, since typically any abstract design can be implemented in many different ways, and changes in implementation may not affect the original requirements, but may alter meta-requirements,
  e.g. robustness or possibilities for deeper integration, or for future changes (Sloman, 2007a);
- developmental transitions in individuals – changes in requirements, designs, and implementations during development – supported indirectly by the genome.
  (Requirements for infants, toddlers, teenagers, parents, etc. all differ.)

There are tradeoffs between processes of evolution and processes of individual development.
The transitions in both may be closely related, as we'll see.
Pressure for them comes partly from the environment, partly from previously evolved forms of embodiment.

Studying interactions between transitions, with increasingly complex feedback loops, may explain some drivers of evolution and development, leading to a better understanding of the products of both.

Specifying what needs to be explained

(What an engineer would call requirements analysis)

A great deal of work by Piaget (and many studying development of cognition) involves experiments to discover what needs to be explained.

Popper posed a problem about doing this: **you cannot just observe – without some theory providing concepts and questions.**

(Popper, 1934). For an introduction see (Magee, 1985)

How observations are described will depend on what the background theory is.

It follows that as scientific theories develop, including producing new ontologies, we often find better ways of presenting what was previously discovered – and then, what needs to be explained is changed.

Piaget's theories and his descriptions of experimental results make use of many concepts that are quite difficult to understand

E.g. accommodation, assimilation – two of the easiest ones), equilibration, conservation, reversibility, integration, horizontal and vertical decalage, pre-operational, concrete operational, formal operational, structure, ...

I believe Piaget's good ideas can be re-interpreted in the light of the idea presented below that intelligent animals find and investigate **often in parallel** many distinct but combinable "exploration domains" and that their learning has several phases including

(i) finding re-usable patterns (associative learning), and then later
(ii) reorganising the material into a generative (proto-axiomatic) system that allows a learner to work out what can happen or must happen in novel situations. (A possibility ignored in (Gibson & Pick, 2000)).

The complexity of the developmental processes

Piaget emphasised that there is no fixed linear developmental trajectory

Perhaps because competences in different domains develop in parallel and interact?

Examples of exploration domains that can interact include:

- Learning about aspects of topology and and aspects of geometry can interact: e.g. continuous geometric changes can produce discontinuous topological changes.
- learning about a new type of stuff and how its properties interact with shapes, or other materials
- learning about a new class of actions that can be performed, e.g. sucking, chewing, grasping (with one hand, with two hands, with teeth, with feet), twisting, dropping, throwing, pushing, pulling, ...

  and many more...

Individual developmental trajectories depend on which domains are being explored, what has been learnt in different domains at particular times, which buggy learning has occurred, which domains are combined – all that can vary enormously between learners.

Many partial successes use inadequate solutions – that don't generalise.

Many special cases develop independently before unifying abstractions are discovered.

Applications of what has been learnt in specific situations can generate multiple types of error, often capable of driving new learning (debugging: (Sussman, 1975)).

A very important type of development depends on the fact that many of the domains involve structures and processes in space and time: a fact that allows domains to be combined to generate new, richer, domains – in very many different ways.

(Contrast learning different branches of algebra, or different programming languages.)

Getting “behind” Piaget’s work.

I’ll start with the phenomena I think Piaget was trying to explain– since I have been exploring many of the same problems,

e.g. (Sloman, 2008b).

I’ll present a few examples of domains (“exploration domains”) that a thinker (young or old) can explore to find hidden structure, where some of the generalisations may be discovered empirically, then later understood as having a kind of necessity.

Some of the examples come from Piaget – others not, though all have features appropriate for use in Piagetian experiments.

I don’t know whether I am merely repeating suggestions already made by Piaget or some of his followers.

Karmiloff-Smith (1992) suggests that Piaget’s theories downplayed the contribution of the environment in evolution, and that he was committed to a uniform (innate) developmental mechanism for all domains even though the details of each domain are different. I am not sure this is an accurate characterisation of Piaget’s position, though I may be hallucinating my own ideas onto his work.

Features of exploration domains

The "exploration domains" described later are of different kinds, but all are domains of activity that can be based on features of the environment

(though some can also be explored in an abstract manner detached from the environment).

- play/experimental/exploration revealing types of occurrence, and correlation patterns.
- discovery of abstractions (e.g. "conservation laws" or "invariants") across examples;
- later simplification by (usually unconscious) invention of new principles underlying patterns discovered:
  1. Finding/creating a more powerful ontology specifying what kinds of states, entities, events, processes and causal interactions can be thought about, perceived, predicted, hypothesised,
  2. Extending the set of kinds of matter (physical stuff) that can occur in those events, processes, etc.;
  3. Discovering new constraints on possibilities and necessary connections between possibilities;
- Combining two independently explored domains to form a richer domain,

  for example learning to play with mud, learning to play with wooden blocks, then learning to play with the combination. (Meccano involves several different sub-domains.)

In some cases there is a meta-cognitive transition that is a precondition for taking part in Piaget's experiments:

- development of meta-cognitive abilities to think about, discuss and answer questions about the things learnt in an exploration domain.

  (Compare Karmiloff-Smith's "representational redescription" of "microdomains")

  Can any non-humans can do anything similar?

Exploration Domains are of different kinds

Here are some of the ways in which exploration domains can differ:

- Some contain discrete structures and processes, others a continuum of possibilities.
- Some involve human language, others non-verbal forms of representation.
- Some of the discoveries are about properties of particular kinds of matter or particular material objects, while others are about more abstract entities relevant to many very different concrete physical situations. e.g. Euclidean lines, areas or volumes.
- Many of the examples involve causal connections between events, processes and structures, where
  - some are deterministic (e.g. moving one end of rigid centrally pivoted rod down causes the other end to go up, and vice versa.)
  - others are statistical, e.g. the proportion of ways of getting a total of 5 when throwing two dice.

This is not an exhaustive classification.

I shall start with a relatively simple discrete, domain which leads to an abstraction that is applicable to many different sorts of things that can be ordered.

A wide range of examples can be found in Piaget's publications and also in Sauvy and Sauvy (1974).
A domain relevant to gaining insight into negative integers is described in (Liebeck, 1990).

Computer programming is an ideal way of exploring a wide range of domains by modelling or simulating the structures and processes.

"Turtle programming" in LOGO was one of the earliest examples illustrating this.

Several more examples of computational exploration of non-computational domains are illustrated in
`http://www.cs.bham.ac.uk/research/projects/poplog/examples`

Explorations of (discrete) possibilities: Order

This exploration domain can use physical objects (cards or blocks with letters) or can be treated as a collection of thought experiments:

A B C D E

If you can repeatedly swap any two adjacent elements you can change the order of letters, e.g. to this using one swap:

B A C D E

Can the next transition be achieved by combining adjacent swaps? How?

A B C D E —> E B C D A

Can you go in the reverse direction using adjacent swaps – and if so, how?

How many different sequences of states (“routes”) are there between the two states with “A” and “E” swapped? How do you know?

Could there be infinitely many different “swap-routes” between two states?

Is it possible to reverse the whole array using a sequence of swaps?

Order transformations and necessity

Consider the previously presented “swap game”, in which cards with letters are presented in a certain order and a target ordering is presented.

The game is played by trying to get from the initial state to the target state by a sequence of pairwise swaps.

A child playing the game may or may not notice various things (and ask “why?”):

- Any card can be moved to the right end by a sequence of swaps, leaving the other cards in order.
- Similarly any card can be moved to the left end.
- The cards at both ends can be made to change position.
- Any two cards can be made to change position.
- Given a row of cards, and another row of similar cards in a different order, the first can be transformed into the second.

  Intermediate stages of comprehension: learning this for specific pairs of rows, then learning it for any two rows of 5 cards, two rows of 6 cards, two rows with any number, etc.
- The procedures work even if some cards in the rows are repeated: e.g. “A B B C” can be transformed to “C B A B”
- The transformation is impossible if one row has more occurrences of a certain letter than the other (e.g “A B B C” “A B C C”
- The colour and size of the letters, the material of which the cards are made (and many more features of the situation) are not relevant to whether the game can succeed or not. (Why not?)
- Would what is learnt in this game be relevant to a game in which there is a collection of nested 2-D figures, e.g. a square inside a triangle inside a circle inside a pentagon?
  What would the “swap” operation become?
  Which of the generalisations true of card-swapping could be applied to the nested figures?
  Why is it a mistake to assume this necessarily involves use of metaphorical reasoning?

A Possible AI project

A possibility for meta-exploration could be undertaken by someone interested in child development who also has AI programming experience.

That could be done by exploring different ways to program a child-like simulated explorer that is capable of:

- perceiving and describing letter sequences
- performing the swap operation on any two items in the sequence
- combining swap sequences, e.g. by writing "virtual programs" to swap virtual cards, and then running the programs.
- noticing patterns that occur when such programs run
- noticing that some patterns never occur

  (e.g. the initial state is never restored after an odd number of swaps)
- reasoning about why things happen or why things do not happen
- examining processes and realising that they have features that are general to a class of processes
- noticing structural mappings between structures, between processes, and between whole domains.

A psychologist exploring such a simulation might notice obstacles to learning, and patterns in how learning occurs in this domain.

This could inspire research into (a) how well some or all children fit such a model and (b) whether brains have mechanisms that could be used for these purposes and (c) how many other animals have such competences and meta-competences, and (d) how such competences evolved, and how many times.

More examples of exploration domains

The examples given so far, involving rows of cards and nested figures, with a “swap” operation are discrete exploration domains: At the lowest level there is a discrete set of entities whose relationships can be changed when processes occur. But there are other sorts.

- People with mathematical minds will quickly realise that the domains presented so far are very much richer than I have indicated, and can be used to explore arithmetical notions, permutations, and other structures.
- Other domains are not discrete but continuous, permitting gradual, and even indiscernible, transitions between states

  (like the “invisible” movement of the hour hand of a mechanical clock).
- Some are a mixture of continuous and discrete – like the domain of movements and configurations of a fixed, discrete set of blocks, meccano pieces, or Lego blocks.
- Some are linear (one-dimensional – totally ordered), some two-dimensional (e.g. embedded in a plane, or in the surface of a sphere), some three or higher dimensional.

  Mathematicians have happily explored infinite dimensional domains.
- Some domains are fixed, static structures whose properties can be explored, while other domains of exploration include events and processes, which themselves may be either discrete, or continuous, or a mixture (think of various kinds of board game).
- Some domains include different kinds of matter – e.g varying in hardness, rigidity, elasticity, fluidity, stickiness, absorbency, solubility in water, electrical or thermal conductivity, and many more – children explore them long before they study physics at school.

Naturally occurring exploration domains

The swap domain may or may not turn up in the life of a child: I chose it as an example likely not to have been encountered by most people.

Anyone who has studied sorting algorithms in computer science, will instantly recognise it, however.
`http://en.wikipedia.org/wiki/Bubble_sort`

An exercise: try to find or think up, naturally occurring exploration domains that can turn up in the life of humans or other animals.

Consider exploration domains relevant to different nest-building birds, for example?

- Nests made by pushing plant matter and perhaps mud or sand along the ground.
- Nests made by bringing discrete objects, e.g. pebbles twigs, which have to be placed to form a rigid structure.
- Nests made with an outer casing and a softer internal lining.
- Nests made by repeatedly bringing a malleable substance such as mud, portions of which can be pushed together, later forming a rigid structure, after drying.

Are there identifiable separate kinds of expertise required for sub-domains that have to be learnt separately, within a nest-building domain?

Contrast different exploration domains in nests on the ground, nests resting on branches in trees, nests suspended from branches, water-edge nests.

Some species of animals (often called “precocial”) seem to be born, or hatched, with all the expertise they need, while others (often called “altricial”) have to build their own competences.

How can a genome encode the information required for either of these?
(Sloman & Chappell, 2005; Chappell & Sloman, 2007)

Compare food-obtaining, and food-consuming domains, for herbivores, carnivores, multivores.

Explorations in a continuous domain: Geometry

There many continuous exploration domains with very different properties.
This is a simple example, related to some of Piaget's examples.

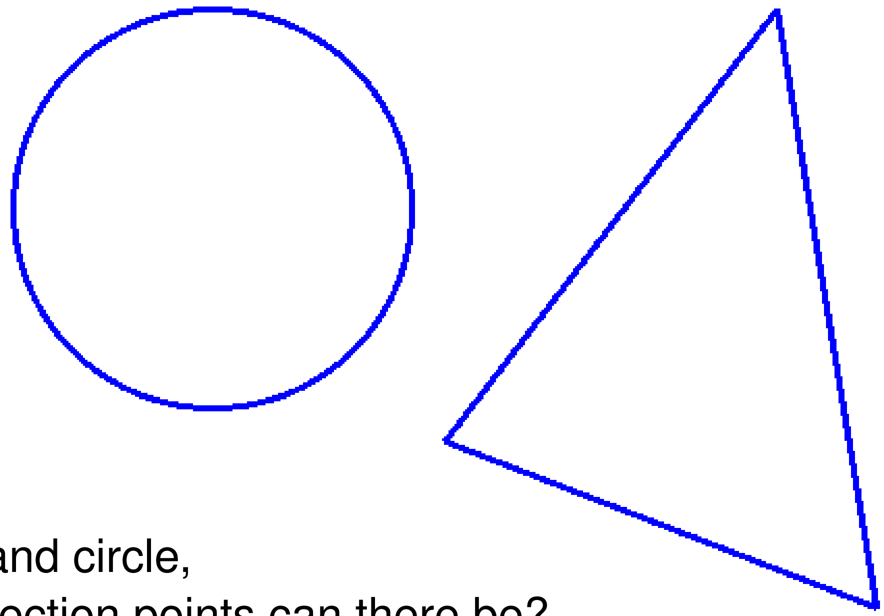

Where else could the circle and triangle have been placed on the surface you see?

In how many different locations could the circle have been placed?

Can locations be counted?

What would happen if the circle were to slide to the right, passing through the triangle?

If you alter positions and sizes of the triangle and circle,
how many different possible numbers of intersection points can there be?

Could there be nine points of intersection? If not why not?

Does the answer depend on the shape of the triangle?

What if the triangle were replaced by a square, or the circle by an ellipse?

Would the answers change if the lines had a different colour?

What difference could the thickness of the lines make?

Continuous domains can be discretised in different ways, for different purposes.

Exploration is not infallible

Not only is it possible to make discoveries in an exploration domain: it is also possible to make mistakes by over-generalising on the basis of an observed pattern.

- For example, someone who tries a limited range of experiments in the swap game may wrongly conclude that if more than six swaps occur in a row then an object must change direction of movement, then later discover that that is false, if the swaps are applied to different objects.
- Likewise, someone who considers only a circle that is very much larger than the triangle may wrongly conclude that it is impossible for a triangle and circle to have more than two, or more than three points of intersection.
- Someone playing with the triangle and circle who moves them only in quite large steps may never notice the case where a vertex of the triangle lies on the circumference of the circle, so that two sides share a point of intersection with the circle,
- It is also possible to miss the case where one or more of the triangle sides meets the circle at a tangent, with only one point of intersection.

A failure to encounter certain possibilities, because the exploration does not include enough generalisation can also prevent the learner acquiring certain concepts, such as the concept of a line being a tangent to a circle, or the concept of a line ending on the edge of the circle.

How many of the errors and confusions shown by children in Piaget's experiments can be explained by models of simulated learners that do not explore enough cases, or form incorrect generalisations? (Compare Seely-Brown's BUGGY program, modelling and diagnosing children's arithmetic errors.)

Exploring continuous deformation of triangles

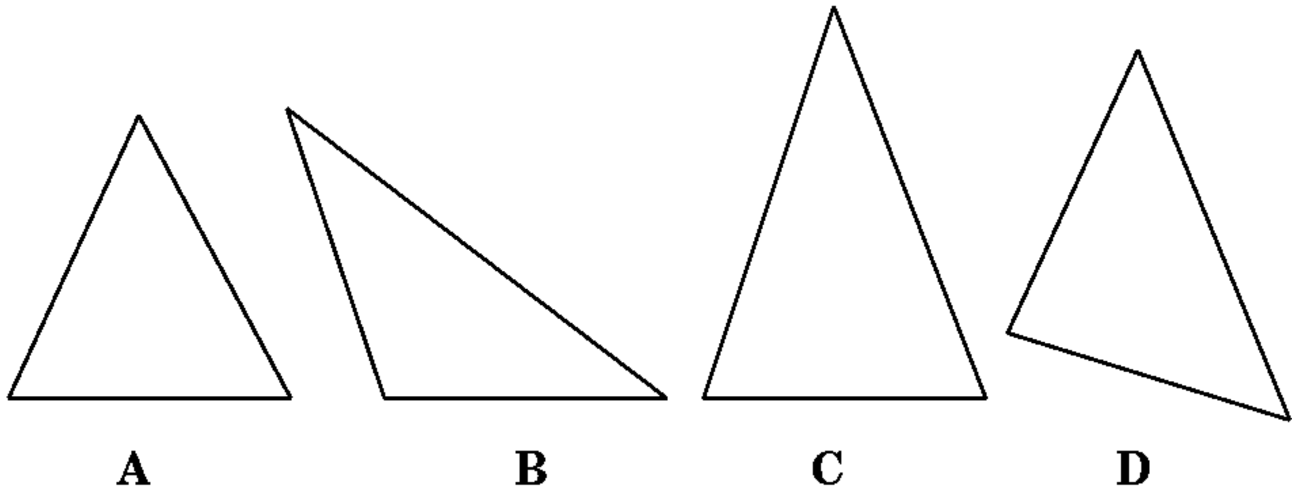

If the triangles are made of rubber bands, held in place by pins what happens to the angles if you move the corners? E.g.

- If you slide the top vertex left or right, parallel to the base, as in **B**?
- If you start with **B** and move the top vertex downwards?
- If you move the top vertex further from the base, as in **C**?
- If you rotate the triangle and deform it as in **D** ?

When lengths of lines and sizes of angles change

- Can all the lengths increase at the same time?
- Can all the angles increase at the same time?
- Can one angle/length change while the other two remain fixed?

If not why not? How could that be discovered?

Explorations of possibilities: blanket, string and toy

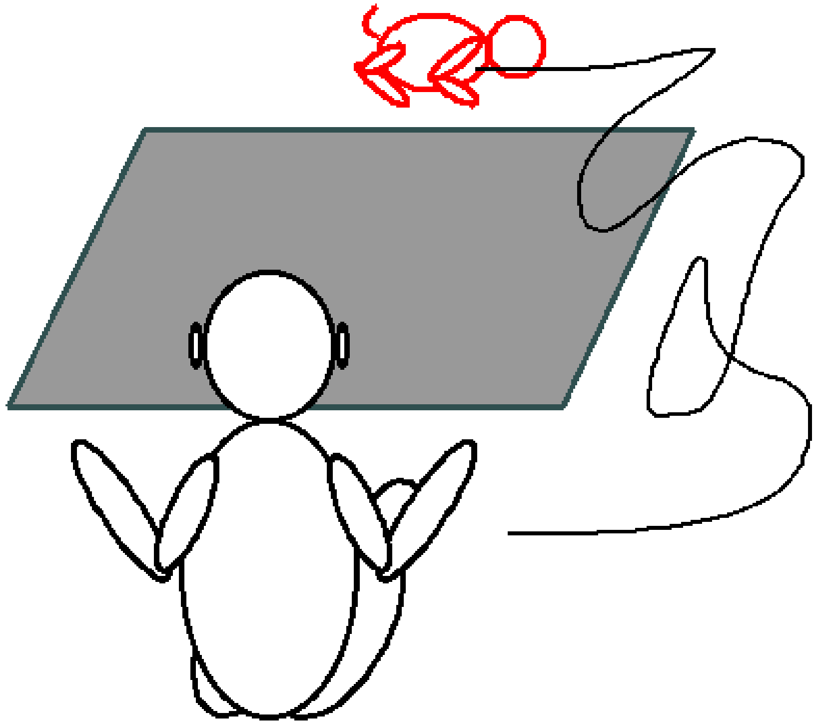

A child, a blanket, a toy
and a string attached to the toy.

This is a domain in which properties of various kinds of matter can make a difference.

- What actions are available to the child?
- What will their consequences be?
- Which ones can bring the toy to the child?
- Which ones can bring the child to the toy?
- What difference does it make if the string is very long?
- What difference does it make if the string is looped round a chair-leg?
- What difference does it make if the string is stretchable?
- What difference does it make if the string is replaced by stiff wire?
- What difference does it make if the toy is on the blanket?
- Or the blanket is replaced by a wooden board with the toy on it?

One of Piaget's examples: Rotating planes

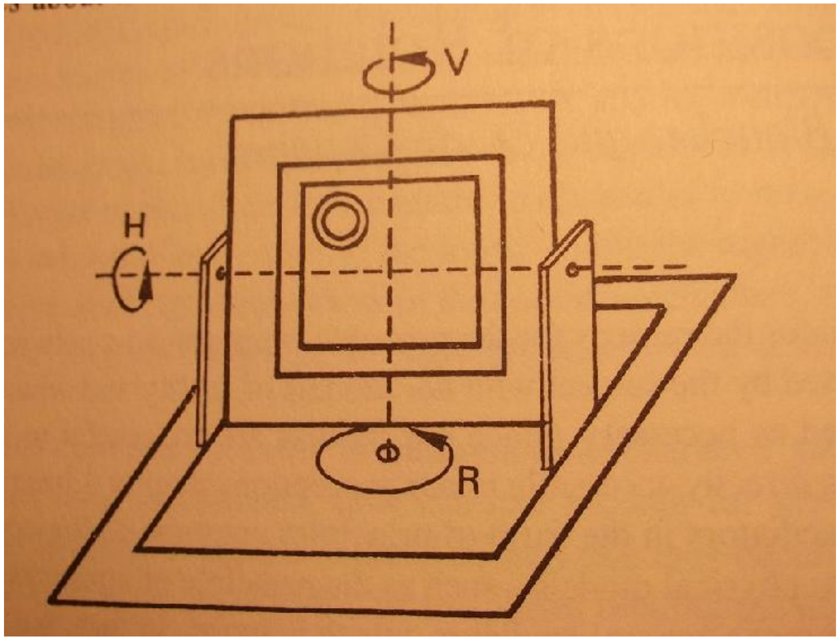

- The central square can rotate about the axis mounted vertically in the outer frame.
- The outer frame can rotate about the axis mounted horizontally across the two pillars.
- The pillars are on a horizontal mount that can rotate about a vertical axis.
- A disc is stuck in the top left quadrant of the central square.

What are the interesting questions that can be asked about

- What processes are possible?
- What states are impossible to achieve from this starting configuration? Why?

Notice that despite the continuous motions this can be treated as a discrete domain: why?

An unexplored domain

I don’t know if any child, or teacher has explored the following issues that arise with a set of cubes, or any other collection of items that can be arranged in different configurations on a flat surface.

- Given a collection of blocks, no matter how many there are, they can be arranged to form a straight line – though in some cases walls, furniture or other objects may limit the number that can be included in the line.
- It is also possible to arrange the blocks in several adjacent lines (rows or columns), where adjoining blocks are lined up.

  o o o o o o o o o o

  can be rearranged as

  o o o o o o o
o o o

  or as two lined-up rows:

  o o o o o
o o o o o

  Whereas this cannot be split into two such lined up rows:

  o o o o o o o o o

  But can be split into three lined up rows:

  o o o
o o o
o o o

  A learner who notices this can discover that there are some rows that cannot be split into similar lined up rows, some that can be split into two, some into three, some into four, etc.

  What else could be discovered about regular arrangements of blocks? What difference does it make if the child cannot yet count?

Investigate your own domains

A good way to understand these ideas is giving yourself the experience of explicitly investigating and then trying to systematize a domain, in contrast with doing it implicitly, as in childhood, driven largely by the environment plus genome.

Piaget seems to have noticed that the world is full of “explorable domains”, or “micro-worlds” that one can play with (or play in) then reflect on and build a theory about.

What can you do with one rubber band and a collection of pins and a pin board to hold a stretched band in place. What becomes possible if you add another band, or two, or three?

What can you do with a flat two-sided object on a table top – e.g. a coin or card with differently coloured faces? What else can you do if there are two such objects?
What happens if you add a third, or a fourth ... ?

If you have coins on a chess board what sorts of patterns and pattern-changes can you generate, using different constraints? (e.g. one coin per square, coins only on black squares, only allow vertical and horizontal, or only allow diagonal changes of arrangements.)

What can you do with a long piece of string?

I have been trying to explore a way of thinking about movable (by translating and/or rotation) line-segments in 2-D space to see how much of Euclidean geometry can be reconstructed without making use of the Euclid’s parallel axiom. This is a rather complex exploration domain. See
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/p-geometry.html`

What can you do with a matchbox, a toothpick and a button?
(I have no idea: I have never tried or thought about it, though as I type this, ideas grow themselves...)

Conjecture: Children (not only humans) spontaneously, or with external encouragement, discover many such exploration domains, and return to some of them repeatedly over weeks, months or years, interleaving and sometimes combining the explorations.

Piaget's experiments on exploration domains

Piaget considered different domains, with different properties, involving

- abstract shapes represented diagrammatically
- physical objects that could be moved (with/without physical constraints)
- liquids and solids
- static structures compared

  Shapes, spatial/topological relationships,
  metrical properties (same or greater length, volume, amount of stuff) etc.
- consequences and limitations of possible changes in a situation
- things changing to achieve a desired result
- questions about what does and does not change (lengths, volumes, amounts of stuff, etc.)
- effects of slope, height and friction on influence of gravity
- Unconstrained consideration of possibilities (what can be done)
- Constrained consideration of possibilities (how to achieve X)
- Constraints from the world: what is impossible
  - Constrained by physics: properties of matter
  - Constrained by geometry: properties of space
  - Constrained by space-time continuity: possible and impossible processes
  - Constrained by topological relations: order, overlap, containment, etc.

I do not know whether Piaget ever described any of his sample domains with mathematical precision, or produced mathematical or computational analyses of their similarities and differences.

A problem with Piaget's experiments

It appears always to be assumed by the experimenters that the children understand the questions they are asked – even if they don't know the answers or don't know how to evaluate and compare answers.

But suppose the process of learning about a new domain is not just

discovering individual facts, and noticing generalisations linking observed facts,

but also

extending the concepts available for formulating questions, factual descriptions, hypotheses, predictions, explanations, desires, intentions, plans, etc., (Sloman, 1978, Ch 2).

Then we need to distinguish (a) experiments probing what the children know from (b) experiments probing what they understand.

However, there are many different ways in which conceptual competences change; and so far we don't have very good theories about that. E.g., there are differences between

- acquiring a new bottom level sub-category in a taxonomy
- adding a new high level branch in a taxonomy
- learning about new relationships (as opposed to categories)
- learning about new kinds of mappings between complex structures (or domains)
- learning new forms of representation supporting concepts and conceptual composition
- learning about new dispositional concepts (fragility, conductance)

  In many cases the new concepts cannot be defined explicitly in terms of the old ones
  (as in scientific advances). (Sloman, 2007b)

Our understanding of mechanisms for supporting such conceptual changes is still very primitive, as shown by the very limited forms of learning in current robots.

Results of the experiments

Unlike much research on children, Piaget's' work does not emphasise statistics of successes and failures, reaction times, times to solve a problem, ages at which various competences are achieved etc.

Statistical presentations (e.g. using averages) can hide important similarities and differences between individual learners, and individual performances of the same learner.

Ignoring those details for the sake of achieving statistical significance is not good science.
(It may sometimes be useful engineering, or medical research.)

In contrast, Piaget and his collaborators probe each child – starting with a standard question or questions and then creatively producing new questions and suggestions on the basis of answers to previous questions.

What their work reveals is that through the whole age range investigated there is wide range of types of

- partial understanding
- misunderstanding
- inability to generalise from specific cases
- over-generalisation from specific cases
- regression from correct to incorrect answers
- use of fallacious reasoning
- construction of bizarre explanations
- inability to combine two or more pieces of information, apparently grasped well
- increases in insight and sophistication over time, though details are different for different children

The results also raise doubts about communication between researchers and children.

Piaget's conclusions

The books provide many interesting insights into the different sorts of thinking, reasoning, explaining that are required for the various tasks presented, but it is not clear to me that the collection of descriptive and explanatory constructs used by Piaget is adequate, even for his own ultimate explanatory purposes.

Instead he needs a collection of more detailed specifications of changing information-processing architectures and mechanisms, that could be used in a learner's mind, employing different forms of representation, different ontologies, different algorithms, and different collections of stored information at different times.

The experimental results (and in some cases school examination results), show that the learning is a long, slow, multi-faceted, often idiosyncratic, and usually also erratic process with large variations in results achieved at various ages, or even in a lifetime.

This could lead some to assume that the biological learning mechanisms are deeply flawed and should not be used in machines.

However that ignores the fact that there are there are changes between birth (or conception!) and adult life in humans and many other species that current machine systems do not even begin to match, though in some domains (e.g. some kinds of data-mining) they far outstrip human learners.

Trying to find out what infants can and cannot do at birth (e.g. looking for "core concepts") is of relatively little interest compared with understanding the processes that produce those various competences (none available at conception!) and continue to develop new ones after birth. (Chappell & Sloman, 2007)

I suspect current theories of learning barely scratch the surface – and fail to specify mechanisms that can account for the changing competences Piaget found, or even the well known patterns of development generally seen in infants and children, and throughout life. So let's not write-off biology.

Types of situation involving possibility and necessity

We could try to specify in far more detail than ever before (but inspired by Piaget) the many different kinds of competence and knowledge that a typical human acquires, and some of the observed developmental and learning sequences, e.g. learning about:

- Discrete variation (externally, internally imposed boundaries)
- Continuous variation
- Linear structures, 2-D structures, 3-D structures, and spatio-temporal structures
- Properties of materials relevant/irrelevant.
- Only kinematics relevant: or kinematics + dynamics.
- Various kinds of change:
  - translations
  - rotations
  - topological changes
  - shape preserving non-preserving
  - breaking/separating
  - joining/assembling
  - changing information (e.g. visibility, occlusion, rotation)
  - caused by: external processes, viewer motion (translation, rotation)
  - caused by intentional manipulative actions/interactions

    actions by self   actions by others (SHOW ROTATING CUBE)

    effects of inanimate objects

    interactions involving single/multiple contact points

  and many more.

  There are many thousands of things for toddlers to learn.

Questions about Piaget

I do not know the answers to these questions:

Did Piaget ever list systematically, in a generative form, how the various problems he considered within a domain could vary, and why.

(My previous slides merely scratch the surface of such a survey.)

I have the impression he thought about collections of examples and features of the collections, but did not attempt to characterise them as a mathematician (or programmer) might specify the sets in a systematic (generative, axiomatic, programmatic) fashion.

I.e. he did not do in his explicit work, what he must have done as a toddler restructuring vast amounts of piecemeal information into systematically applicable micro-theories – as conjectured below.

Did he ever discuss (explicitly) information processing mechanisms, forms of representation, ontologies, architectures, ... or how these change?

(All of these central to AI/Robotics)

He started moving in the right direction (towards explanatory mechanisms) in considering “conservation” and things that could be expressed using various mathematical ideas (commutativity, associativity, abstraction, invariants).

But did he have a clear vision of where that could lead, e.g. in AI?

Late in life he said he wished he had known about AI as a young researcher
(at a workshop on Genetic Epistemology and AI, in Geneva, circa 1983).

I’ll now move on to a conjecture about how knowledge gained about an exploration domain can change.
(Compare this with what Annette Karmiloff-Smith calls representational re-description?)

Embodiment

In recent years the belief has spread that some older approaches to the study of cognition (e.g. using symbolic forms of representation, such as logic and algorithmic knowledge) are erroneous, because cognition somehow rests on and grows out of features of embodiment including the interactions between an agent's body and the environment

For example, it was found many years ago that giving robots compliant wrists can simplify some manipulation tasks by allowing some of the details of wrist motion to be controlled by interactions between surfaces of objects rather than having to be precisely controlled centrally – e.g. when grasping an object the precise motion of the hand can be controlled directly by the surfaces of the object, rather than by a planning and plan execution sub-system that uses information about the shape of the object.

However, this emphasis on embodiment ignores the fact that there can be structures in the environment that are perceivable and manipulable by animals or robots that have very different morphologies and sensorimotor mechanisms (e.g. birds, primates, elephants).

So biological evolution may well have found a use for methods of learning and solving problems that abstract away from the details of sensorimotor morphology and the detailed dynamics of physical interaction.

This could explain how humans born with various physical deformities, including missing limbs, malfunctioning control systems (palsy), blindness, etc. can develop a rich understanding of the world shared with others.

The strong embodiment thesis also ignores cases of off-line interaction with the environment, in predicting, explaining, planning, describing, narrating, questioning, etc.

I think it is clear that Piaget does not accept the "strong embodiment" thesis.
That's the right response to it! (Sloman, 2009)

The basis of mathematical knowledge

Kenneth Craik suggested that some animals have abilities to “work out” consequences of and constraints on possible actions instead of always having to learn by trial and error (Craik, 1943).

That can be very important when error means death or serious injury – or even just a huge waste of time and energy.

This can be seen as another version of the conjecture that some of the main mechanisms that produce “U-shaped” language learning (starting pattern-based, then develop a more systematic, generative theory usable in novel situations) are involved in forms of learning about domains of exploration, in humans (starting before they learn to talk) and other species.

For more on this see the ideas here about “toddler theorems” and the evolution and development of mathematical competences:

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler`

Intelligence, especially mathematical intelligence, largely consists in productive laziness.

Are humans special? How

I suspect that in humans and many other species this form of learning by exploring then reorganising to form a productive system (like going from pattern-based to syntax-based language use) happens in connection with different domains at different times.

However it seems that only humans have the meta-cognitive mechanisms required to explicitly notice, think about, and talk about the kinds of learning they (and the others around them) experience – meta-domains of exploration can also have a developmental trajectory.

A consequence of being such a species is that the use of cultural transmission to accelerate child development can achieve more and more with each successive generation.

But there must be a limit and we may be close to it: the singularity of cognitive catchup, described in
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/another-singularity.html`

Some questions to be discussed later

1. Can the postulated development from pattern-based to generative information about a domain also be achieved by evolution, or can it happen only during individual development?

2. If it can happen in evolution does that explain some examples of creative problem solving in other animals – using inherited competences?

   (Most animals, e.g. microbes, insects, ... seem to have only inherited competences.)

3. There's lots to be said about how useful, increasingly high level, ontologies can be compiled into the perceptual and action mechanisms of an organism or machine, either by evolution or by individual learning and development.

   (This is a requirement for many innate competences)

   (There are deep implications for the study of vision, for instance.)

4. to be continued ...

To Be Continued

This is a draft set of slides, to be revised and extended as new ideas, or criticisms turn up.

In particular I shall later try to add a section on why it seems likely that young children and many animals do not use a global 3-d cartesian coordinate system to represent the environment and its contents, but instead use a dynamically changing web of **partial** orderings of directions distances, sizes, slopes, curvatures, velocities, accelerations, forces, etc.

Darke (1982) discusses and criticises “The Topological Primacy Thesis”, attributed to Piaget and Inhelder. I don’t think that what I am presenting is the same thesis even though I believe that the sensory and motor resources of biological organisms initially do not have access to a metrically based global coordinate system, and instead have to make do with a complex (and changing) web of partial orderings and topological relationships (including contact, overlap, containment, discontinuity, etc.) The Piagetian experiments described by Darke, are not relevant to testing my hypothesis: it would have to be tested by building and running a working model of the perceptual, learning, reasoning and acting system based on partial orderings, though how to do that is still far from clear.

There’s much more to be said about all this; and much more to be done by roboticists.

I think there are important connections with Karmiloff-Smith’s work on “Relational Redescription”, though I have not yet tried to compare in any detail. Karmiloff-Smith (1992, 1994)

One difference seems to be that K-S specifies that RR involves consciousness, whereas I distinguish the conceptual revision (e.g. from pattern based to syntax-based forms of representation) from the meta-cognitive reflective step of becoming aware of what has been learnt.

I have been concerned with how these processes provide a basis for human mathematics, which links up with Kant’s and Piaget’s ideas about necessity, and with Craik’s theory about use of models.

Luc Beaudoin has pointed me at work by Carl Bereiter, that may be relevant.

References

Borovik, A. (2010). *Shadows of the Truth: Metamathematics of Elementary Mathematics*. Mancester University: Unpublished Draft.

Chappell, J., & Sloman, A. (2007). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, *3*(3), 211–239. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609)

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

Darke, I. (1982, May). A Review of Research Related to the Topological Primacy Thesis. *Educational Studies in Mathematics*, *13*(2), 119–142.

Gibson, E. J., & Pick, A. D. (2000). *An Ecological Approach to Perceptual Learning and Development*. New York: Oxford University Press.

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

Karmiloff-Smith, A. (1994). Precis of: Beyond Modularity: a Developmental Perspective on Cognitive Science. *Behavioral and Brain Sciences*, *17*(4), 693-706.

Liebeck, P. (1990, Jun). Scores and Forfeits: An Intuitive Model for Integer Arithmetic. *Educational Studies in Mathematics*, *21*(3), 221–239.

Magee, B. (1985). *Popper*. London: Fontana (Modern Masters Series).

McCarthy, J. (2008). The well-designed child. *Artificial Intelligence*, *172*(18), 2003-2014.

Popper, K. (1934). *The logic of scientific discovery*. London: Routledge.

Rochat, P. (2001). *The Infant's World*. Cambridge, MA: Harvard University Press.

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

Sloman, A. (1978). *The computer revolution in philosophy*. Hassocks, Sussex: Harvester Press (and Humanities Press).

Sloman, A. (2007a). *A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots* (No. COSY-DP-0701). (Contribution to euCognition wiki)

Sloman, A. (2007b). *Why symbol-grounding is both impossible and unnecessary, and why theory-tethering is more powerful anyway.* (Research Note No. COSY-PR-0705). Birmingham, UK. (http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models)

Sloman, A. (2008a). *Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)?* (Research Note No. COSY-PR-0702). Birmingham, UK. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0702)

Sloman, A. (2008b). *A New Approach to Philosophy of Mathematics: Design a young explorer, able to discover "toddler theorems".* (Online presentation http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler)

Sloman, A. (2009). Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. In B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, & K. Doya (Eds.), *Creating Brain-like Intelligence* (pp. 248–277). Berlin: Springer-Verlag.

Sloman, A. (2010). If Learning Maths Requires a Teacher, Where did the First Teachers Come From? In A. Pease, M. Guhe, & A. Smaill (Eds.), *Proc. Int. Symp. on Mathematical Practice and Cognition, AISB 2010 Convention* (pp. 30–39). De Montfort University, Leicester.

Sloman, A., & Chappell, J. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings IJCAI'05* (pp. 1187–1192). Edinburgh: IJCAI. (http://www.cs.bham.ac.uk/research/cogaff/05.html#200502)

Sussman, G. (1975). *A computational model of skill acquisition*. San Francisco, CA: American Elsevier. (http://dspace.mit.edu/handle/1721.1/6894)

Tarsitano, M. (2006, December). Route selection by a jumping spider (Portia labiata) during the locomotory phase of a detour. *Animal Behaviour*, *72, Issue 6*, 1437–1442. (http://dx.doi.org/10.1016/j.anbehav.2006.05.007)

Turing, A. (1950). Computing machinery and intelligence. *Mind*, *59*, 433–460. ((reprinted in E.A. Feigenbaum and J. Feldman (eds) *Computers and Thought* McGraw-Hill, New York, 1963, 11–35))