Understanding causation in robots, animals and children: Hume's way and Kant's way.

Aaron Sloman

`http://www.cs.bham.ac.uk/~axs/`

University of Birmingham, UK

The main topic (ways of representing, learning about, and using information about causation) is in Part 3

These slides also provide some background information about the CoSy project, about some relevant aspects of methodology of science (due to Lakatos), and some speculations about how evolution produced mathematicians, among other things.

NOTE: These slides are accessible from two sites and on slideshare.net
`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#mofm-07`
`http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07`
Also to be installed on slideshare.net `http://www.slideshare.net/asloman/presentations`

Why did I join the CoSy project?

Answer:

Because I wanted to understand how to build a child mathematician: as part of expanding and defending Kant's views on mathematics and causation – which is what originally led me from mathematics into philosophy, and then into investigations of forms of representation, vision, architectures, self-observation, biological origins, cross-species comparisons, robotics ..... a better way to do philosophy than arm-chair theorising.

There are two very different ways of thinking about these topics:

Humean:

All we can think about is what we derive from our own experience (concept empiricism).
In robotics, and psychology, that can lead to the idea that all knowledge, even all thinking, is only about sensorimotor structures and relationships.
I.e. the ontology available is purely **somatic**, though it may be **multi-modal**
and causation is just a matter of experienced correlations.

Kantian:

We have to use concepts that are not derived from experience, in order to have experiences
   (e.g. concepts of spatial relationships and organisation).
Some of our knowledge and thinking is not about sensorimotor structures and relationships but about what is going on in an environment that exists independently of our sensing it.
I.e. the ontology available is partly **exosomatic** and **amodal**
Some causation is about interactions between structures in the world.

Einstein on problems vs. solutions

"The mere formulation of a problem is far more often essential than its solution, which may be merely a matter of mathematical or experimental skills.

To raise new questions, new possibilities, to regard old problems from a new angle requires creative imagination and marks real advances in science."

I've seen that quoted in many places, but nowhere with a reference to the original.

(It is often quoted without "often", as if things could be more or less essential.)

Transforming the slogan for engineers:
Getting the requirements specification right is often the hardest and most important task.

For scientists:
Our ability to construct deep new explanations is limited by our ability to identify what needs to be explained.

For philosophers
Before arguing about necessary and/or sufficient conditions, be sure to ask:
FOR WHAT? That question may not have a unique answer.

Overview

1. **Prologue (why I am doing this)**

2. **Background to the main topic**

   2.a. Some Conjectures about Evolution

   2.b. On progressive vs degenerating research
   programmes (Lakatos)

3. **Two Kinds of Causal Understanding**

   (The main topic)

   With stuff on vision, forms of representation, ontologies, precursors to human
   language, ...

1. Prologue (why I am doing this)

How I got into AI and eventually into the CoSy project.

It all started when I was working on my Philosophy D.Phil thesis (completed 1962) trying to show how Kant's ideas about mathematical knowledge were superior to the prevalent ideas among analytical philosophers at that time, who tended to agree with Hume: all knowledge is either empirical or trivial.

Kant said mathematical knowledge was one of the exceptions, since it was both non-empirical and synthetic (non-analytic, non-trivial), and as a mathematics graduate I knew Kant was right.

Explaining why he was right was another matter: it required showing how making a mathematical discovery expands your knowledge, as opposed to merely making definitions and their implications explicit.
  (E.g. 'All bachelors are unmarried')

It was clear to me that some important ways of doing mathematics were very closely related to our ability

– to see things, e.g. spatial structures,

– to see ways in which they could be transformed

– to see what the consequences of such transformations would be.

  (Later I learnt that this was connected with Gibson's theory of affordances.)

This led me into work on forms of representation and the study of vision in AI, in the 1970s.

That and other things led into architectures, motivation, emotions and other forms of affect.

I fear we still have not formulated the real problems well, because we focus on much easier ones.

I am not an engineer – though I produce software

Why am I in CoSy? Part of the answer:

- Because I became convinced about 35 years ago that the best way to make progress in certain areas of philosophy is to design working fragments of minds

- Because in order to solve problems in the philosophy of mathematics (about the nature of mathematical concepts, knowledge, and proofs) we need to design a working mathematician, e.g. starting as a young child not yet able to count.

- Because human mathematical abilities (concerned with numbers) are side-effects of biological evolution's solutions to some design problems – E.g.
  - How to perform sequences of operations, e.g. stepping, hitting something, reciting lists of names.
  - How to perform two or more related sequences of operations in parallel
  - How to learn from and find new uses for abilities to perform sequences
  - How to observe and control (or debug) sequences of operations performed in parallel (requires at least one more process run in parallel)
  - How to perceive, produce and control processes involving interacting 3-D structures.
  - How to learn about things (and debug your theories) by performing actions.

NB: Studies of 'subitizing' often miss the main points about understanding numbers.

Some of the above ideas are in Chapter 8 of The Computer Revolution in Philosophy (1978):
`http://www.cs.bham.ac.uk/research/projects/cogaff/crp`

CoSy's Aims: New Knowledge, not New Machines

FROM THE COSY WEB SITE (www.cognitivesystems.org):

The CoSy project is a four-year project, involving 7 European Universities. The project is inspired by the visionary EC Framework 6 objective for 'Cognitive Systems':

"To construct physically instantiated ... systems that can perceive, understand ... and interact with their environment, and evolve in order to achieve human-like performance in activities requiring context-(situation and task) specific knowledge"

We assume that this is far beyond the current state of the art and will remain so for many years. However we have devised a set of intermediate targets based on that vision. Achieving these targets will provide a launch pad for further work towards the long term vision.

In particular we aim to advance the science of cognitive systems through a multi-disciplinary investigation of *requirements, design options and trade-offs* for human-like, autonomous, integrated, physical (e.g. robot) systems, including requirements for architectures, for forms of representation, for perceptual mechanisms, for learning, planning, reasoning, motivation, action, communication and self-understanding. The results of the investigation will provide the basis for a succession of increasingly ambitious working robot systems to test and demonstrate the ideas in scenarios that require integration of components that are normally studied separately in various sub-branches of AI, Cognitive Science and Psychology.

How I joined CoSy

In 2003, I heard about the EU Cognitive Systems initiative, mostly driven by Colette Maloney, a physicist working as an EC research administrator. I was amazed: I had been trying for years to get funds for this kind of research and always failed because the ideas were too ambitious, or too vague or challenged too many assumptions held by reviewers.

The EU initiative emphasised that the main goal was to expand our knowledge: building machines was just a means to increased understanding.

So when Henrik Christensen invited me to join a consortium he was assembling it seemed to be an ideal way to continue down the path I had started exploring.

I knew the problems would be very hard and that we would not end up with an impressive working robot, but expected that we would at least learn things that others were not learning because they were looking in other directions.

In particular, our workplan had heavy emphasis on investigating requirements, though not enough people have learnt to do that.

It turned out even harder to make progress on working implementations than I had expected – partly because many AI researchers, especially vision researchers, had abandoned most of the hard problems (e.g. perception and understanding of structure and function, and how to put many capabilities together), and had been working on much easier problems (e.g. learning using statistical techniques).

But relating CoSy to problems in biology, psychology and philosophy, has helped me identify more clearly than before what some of the gaps in our understanding are.

For some suggestions regarding how to identify and fill the gaps, see this summary paper (2 pages):

> http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0608
> How to Put the Pieces of AI Together Again (2006, AAAI Conference, members poster session.)

Part 2

2. Some Conjectures about Evolution

Trying to understand psychology by studying only humans is like trying to understand physics by studying only a type of physical machine made by one successful company.

- Most philosophers tend to think there is one kind of mind and try to specify its necessary and sufficient features.

- Most psychologists studying vision, learning, problem-solving, emotions, language understanding, social interaction, etc. study only humans.
    (There are exceptions, of course, including some at this workshop.)

- If we want deep understanding we need to study the space of possible designs for information-processing systems and the space of requirements for such systems, and how the two spaces relate.

  That includes studying damaged and un-developed designs, and how they relate to various niches.

- This approach leads to some conjectures about evolution and development, and the diversity of products of both, presented below.

- Some of my conjectures came from being puzzled about some of the differences found in nature, e.g. between precocial and altricial species, and the speed and creativity of some kinds of learning in humans and some other animals, compared with other forms of learning.

Development of the ideas accelerated after I started working with a biologist, Jackie Chappell.

Three joint papers, and our workshop slides on causation, are listed at the end, with links.

In defence of speculation

We should not be afraid to speculate, as long as the speculations are part of a research programme in which precision, generality, and the breadth of what is explained and predicted all steadily increase.

See the work of Imre Lakatos on the methodology of scientific research programmes, and his distinction between progressive and degenerating research programmes

– a strong corrective to naive falsifiability theory (often based on mis-reading Popper).

NB:

From this viewpoint, scientific theories cannot always be tested decisively at any particular point in time: rather evaluation can take years or decades.

That's why it is always a mistake to reject a theory just because it is not falsifiable and does not make new predictions: it may be a reformulation of old ideas that eventually becomes far more progressive than the old ideas – e.g. by provoking new sorts of empirical research.

(In my case that includes research on ontology development – examples later.)

Imre Lakatos, (1980).

The methodology of scientific research programmes, Philosophical papers, vol. I,
Eds. J. Worrall and G. Currie, CUP.

This web site has a useful summary of his main ideas: `http://www.answers.com/topic/imre-lakatos`

Some Observations, Conjectures, Speculations

1. Evolution produced far more design solutions for far more problems than we have so far discovered. Identifying those problems is not easy.

2. You can't understand designs without knowing what problems they solved – and a complex design may have solved a very large number of problems.

3. You cannot understand any design without understand how and why it differs from other designs in its neighbourhood in design space (alternative designs) and how it relates to neighbourhoods in niche spaces (alternative sets of requirements).

4. Solutions to problems may have important and useful side-effects for which they were not selected, and not every important feature has a mechanism producing it.

   (E.g. Be very suspicious of a box labelled "emotion" in an architecture diagram.)

5. Human competences DO NOT scale up (Humans are defeated by combinatorics.)

6. Human competences DO scale out – they are combinable with others in creative ways.

7. A model of some competence that works and performs well on its own terms may be incapable of being combined with others as required. (It may not scale out).

8. Many novel combinations of human competences arise in 3-D spatial structures and processes: spatial coincidence and proximity (with and without temporal proximity) supports novel combinations of structures and processes.

9. We do not yet understand what forms of representation of 3-D structures and processes humans and other animals use.

10. A human without a normal body nevertheless inherits mechanisms that evolved in normal human bodies: for human minds to develop, human embodiment is not as important as having had embodied human ancestors. (Or being designed as such.)

Philosophy cannot be avoided

Everyone who is ignorant of philosophy is doomed to reinvent it – badly

(Apologies to Santayana)

A major task of philosophy is to clarify concepts and disputes that arise from conceptual confusions.

Many concepts used by scientists and engineers generate confusions because they become divorced from the ordinary contexts that make them useful, e.g.

emotion, belief, consciousness, autonomy, learning, understanding,

truth, causation, self, reality

Example 1: much recent work on consciousness assumes "consciousness" is a unitary concept, whereas the ordinary notion that we all understand is a mish-mash of concepts of different sorts.

So there is no one thing referred to by the noun, whose functions, evolution, brain mechanisms, (etc.) need to be explained.

Example 2: many debates on which people are tempted to take sides are empty of content, e.g.
- whether abstract entities such as numbers, and proofs really exist
- whether things existed before anyone or anything perceived them
    (as if being perceived were the only kind of effect something can have)
- whether virtual machines and other high level aspects of complex systems are "less real" than the physical implementation machines.

Does it matter if people hold strong views about some vacuous questions?
Yes – it may stop them asking some important research questions.

E.g. about virtual machines in biology.

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#mos09`

Three conjectures

1. Evolution 'discovered' the need for and the power of virtual machines and their role in self-monitoring and self-control, long before we did.

They are used in humans and other animals .

2. Evolution 'discovered' that it could be useful for some animals to treat other organisms (e.g. predators, prey, competitors, conspecifics) as containing internal virtual machines.

This uses the design stance and is more general and more powerful than the intentional stance. It does not presuppose rationality in the observed agents.

3. Evolution 'discovered' ways of improving the search for good solutions by separating specification of high level features from specification of implementation details.

The next slide formulates these three conjectures in more detail.

NOTE: A running virtual machine can do things, and can be effected by things.
It should not be confused with a computer program, which is just a static structure.
(Explained in more detail here   `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#inf`)

Three conjectures about virtual machines and evolution

1. Evolution 'discovered' that for organisms with very complex control systems, the internally directed self-monitoring systems should not just monitor details of chemical/neural mechanisms (as happens in homeostatic systems, for instance) but also the high level virtual machines running on those mechanisms, e.g. machines in which percepts, beliefs, desires, plans, inferences, wishes, puzzles, problem solving, and other things occur and interact causally, producing both virtual machine effects and physical effects.

Socrates: 'Know thyself'. Response: Impossible to do effectively except at the virtual machine level.

2. Evolution 'discovered' that for sophisticated organisms to be able to predict or explain the behaviour of others, and to interact socially with them, it is insufficient to focus only on observable behaviours of others, and impractical to use hypothesised but inaccessible details of neural/chemical processes, but both useful and practical to hypothesise that virtual machines are running in others, exercising high level control.

Dennett (and Newell 'the Knowledge Level') would say: doing that requires taking the intentional stance and assuming others are rational. But for many aspects of information-processing rationality is irrelevant: the virtual machine design stance is generally more productive than the intentional stance. I don't need to take the intentional stance to realise that a toddler is trying to do something that will harm him, nor is it needed to explain how he learns to keep his balance when walking, or learns to talk.

3. Evolution 'discovered' that in order to give sophisticated information processors improved capabilities it was more effective (a) to search in a space of virtual machine designs and to separate that problem from the problem of how to implement the designs than (b) to try to evolve new brain mechanisms directly.

That search may have benefited from separating out various aspects of the genome involved in building something complex so that the different aspects can evolve separately, and especially separating specifications for generic functionality from 'parameters' that instantiate that functionality.
This could include something like separating VM specifications (the parameters) from compilers/interpreters for the specifications (the implementation mechanisms) in the genome.
(Compare work in evolutionary computation using GPs rather than GAs.)

A fourth conjecture: Added 27 Nov 2009

This Kantian conjecture should have been included in the original list:

4. Evolution "discovered" the advantages, for some organisms, of using not only somatic (sensorimotor) ontologies and forms of representation (modal and multi-modal), but also exosomatic ontologies and amodal forms of representation.

This may sound like Kant's "things in themselves" (noumena) but can be expressed as a scientific hypothesis rather than a metaphysical thesis.

The conjecture applies to a subset of organisms (e.g. organisms whose manipulators are not rigidly fixed to eyes, such as nest-building birds and nut/fruit/berry-picking mammals) and those that need to be able to reason about routes across terrain.

Any process in which 3-D objects change relationships (e.g. one surface pushing another past a stationary object) can project into infinitely many different retinal process patterns; and if produced by an animal, that process can be generated by a huge variety of possible sequences of motor signals.

I conjecture that evolution "discovered" how to reduce the combinatorial explosion of patterns of sensorimotor signals, by using exosomatic ontologies with amodal forms of representation, referring to things whose existence does not depend on their being perceived or manipulated.

An example is perceiving a rotating 3-D cube (wire-frame, or solid) and representing the process using a 3-D ontology, instead of constantly trying to predict what will happen next in 2-D retinal stimulation.
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/nature-nurture-cube.html`

Examples of this sort indicate that conscious experience can in some cases include processes occurring in different domains of structure in parallel: e.g. retinal and other sensory processes, and processes in 3-D space in which objects move, rotate, and change multiple relations in parallel (multi-strand processes). See: `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cons09`

Benefits of using amodal exosomatic information

There are very many, very varied, 3-D structures and processes that can occur in the environment, including large subsets that can be produced by individual organisms or robots acting in the environment.

Each such 3-D process can produce a huge variety of sensory patterns depending on viewpoint and motion of perceiver or which body part is used for haptic perception: a doubly-infinite combinatorial explosion of combinations of information fragments.

- Examples include relations between portions of surfaces of objects, such as the opposing surface patches on finger and thumb with an object between them, as required for some forms of grasping.

- By using exosomatic ontologies, and amodal forms of representation, an organism or robot can reason economically about a 3-D process that can be produced in a multitude of ways, plan future actions, think about past actions, predict what will happen, and generalise from things learnt in specific contexts.

- If that sort of representational competence is available, then grasping done by another individual will simply be another special case of two facing surfaces moving together with an object between them.

- "Mirror neurons" may be a special case of something more basic and general that probably evolved earlier ("process abstraction neurons"), because of the need to perceive affordances, plan actions, etc., without representing irrelevant details, and was perhaps later used for perceiving vicarious affordances, e.g. affordances for (and actions of) offspring, predators or prey.

- This still leaves the problem of how such perception works, and exactly what representations are used.

- It requires an ability to 'project' from the exosomatic ontology to a somatic sensory or sensorimotor ontology, and to use sensor information to derive exosomatic percepts (probably in registration with the sensory details) – but such competences seem to be needed anyway for many purposes.

- Robot SLAM systems (Self Localisation and Mapping) demonstrate these features, but currently robot 3-D manipulation using hands is still mostly restricted to somatic, sensorimotor ontologies.

Risks in looking for mechanisms before requirements.

Designers of complex control systems and intelligent robots, need to learn what problems evolution solved, what solutions it produced (e.g. designs, forms of representation, mechanisms, ...) and how they are solutions to those problems.

Otherwise biological designs will not be understood properly, and biologically-inspired designs produced with good intentions will get things badly wrong.

Such research going from problems to mechanisms is very different both from
- going top down from high level designs to low level mechanisms
- working bottom up starting by attempting to emulate the supposed mechanisms of biological evolution and biological learning, which is mostly what 'biologically inspired' researchers do.

Of course, working bottom up is a good way to learn more about the potential and limitations of a class of mechanisms: I am merely countering exhortations to work only that way.

If researchers work only on what can be explained by supposed mechanisms, deliberately ignoring all other considerations, they may never discover that their catalogue of biological mechanisms/designs is incomplete.

Trying to identify the problems that were solved, independently of the mechanisms counters that risk.

(Of course, this warning applies to any group of researchers who assume that all solutions must be built up from some particular class of mechanisms, e.g. Turing machines, von Neumann machines, logical machines, dynamical systems, etc.)

Extending Nature/Nurture options and tradeoffs

Often nature/nurture discussions consider too few options - and proponents of the options don't realise that they are all slinging mud at straw men: the variety and complexity of biological mechanisms is barely understood.

For instance it is often assumed that the only way evolution can avoid a "tabula rasa" and provide information about the environment is to provide a collection of fully formed, innately determined "modules" (a view criticised by Annette Karmiloff-Smith.)

If we take up the design stance we can see many more options, some sketched below, building in part on new ideas about the importance of epigenesis, and the varieties of epigenetic mechanisms.

Compare Eva Jablonka and Marion J. Lamb, 2005,
Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life, MIT Press.
BBS precis: `http://www.bbsonline.org/Preprints/Jablonka-10132006/Referees/`

See the Sloman & Chappell commentary, BBS, 2007 cited at the end of this presentation.

In particular, let's speculate about

- the various levels of abstraction at which evolution can provide 'preconfigured' information about the environment,

- ways in which it can delay construction of competences until relatively late in learning and development

- ways in which development of the ability to learn about the environment might be cascaded
  (as suggested by Oliver Selfridge long ago.)

Human engineers find it useful to design and create implementation-independent virtual machines

Could evolution have done that before us?

- Human engineers have found feedback and feedforward control systems very useful – and evolution produces them all over the place.

- Human engineers find it useful to separate design specifications into a generic part and parameter part: Could evolution 'discover' the benefits of protecting a generic part from mutations while allowing the parameter part to vary much more, and vice versa?

  from something like   GGGXGGGXXGGGGXGGGG.........
  to something like   GGGGGGGGGGGGGGG.....XXXX.....

- That might also allow the generic part to be copied several times during development of the individual and combined with different, separately evolved parameters.

  something like   GGGGGGGGGGGGGGG*3.....XXXX..YYYYY...ZZZZ..
  (Evolving several variants of the same skill – e.g. kinds of walking or running perhaps?)

- That might be followed by inventing a mechanism for deriving parameter parts from interaction with the environment during development of an individual

  (I.e. switching learning about the environment from phylogenesis to ontogenesis).

- Later, some parameter parts might also be segmented into a part that gets built early in the life of the individual and is then used to build larger parameters by getting different sub-parameters from the environment. (A common feature of epigenesis?)

- To see how any of this might work, we need to understand both the environments that pose the problems evolution solves, and the constraints on evolutionary searching.

Two quotations – not contradictory:

Annette Karmiloff-Smith:

"Decades of developmental research were wasted, in my view, because the focus was entirely on lowering the age at which children could perform a task successfully, without concern for how they processed the information."

Karmiloff-Smith, A. (1994).
Precis of Beyond modularity: A developmental perspective on cognitive science.
Behavioral and Brain Sciences 17 (4): 693-745.
(Preprint: `http://www.bbsonline.org/documents/a/00/00/05/33/index.html`)

We could make a similar comment about people studying under what conditions various animals do interesting things, without explaining how they do any of those things.

Ulric Neisser:

"We may have been lavishing too much effort on hypothetical models of the mind and not enough on analyzing the environment that the mind has been shaped to meet."

Neisser, U. (1976)
Cognition and Reality San Francisco: W. H. Freeman.

Comment: We need both

- deep understanding of important PROPERTIES OF ENVIRONMENTS animals or children interact with, and
- deep theories about the INFORMATION PROCESSING MECHANISMS that make it possible to engage fruitfully with those environments

Should we teach students to study environments in greater depth?

John McCarthy on The Well Designed Child

Most AI researchers try to find a small number (preferably only one) of powerful, general, learning mechanisms that can learn from arbitrary data.

Contrast John McCarthy:

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

...

"Animal behavior, including human intelligence, evolved to survive and succeed in this complex, partially observable and very slightly controllable world.
The main features of this world have existed for several billion years and should not have to be learned anew by each person or animal."

Unpublished online paper: 'The well designed child',
`http://www-formal.stanford.edu/jmc/child.html`

Biological facts support McCarthy:

Most animals start life with most of the competences they need – e.g. deer that run with the herd soon after birth. For them, there's no blooming, buzzing confusion (William James)
So why not humans and other primates, hunting mammals, nest building birds? ... and some future robots
Perhaps we have not been asking the right questions about learning.

We need to understand the nature/nurture tradeoffs, much better than we currently do, and that includes understanding what resources, opportunities and selection pressures existed during the evolution of our precursors, and how evolution responded to them.

This requires us to understand the environments involved, as well as mechanisms, architectures, etc..

We need to expand alternatives to be considered

- Evolutionary knowledge about the environment could be partly encoded in strategies for learning about 3-D structures and processes by performing experiments, and for debugging what is learnt.

- What appears to be random 'motor babbling' in an infant could be part of a controlled set of experiments.

- I'll show some examples where there clearly seems to be goal-directed action.

- Humans and animals are not unitary entities so that you can ask: what does it perceive, know, want? There are many different subsystems operating in parallel, and they need not communicate fully.

- It's often taken as obvious that all concepts must come from experience, using a process of abstraction ("concept empiricism", recently re-invented as "symbol-grounding" theory). I've shown that far from being obvious it is false (using arguments from Kant and 20th century philosophers of science).

- Not all new concepts are definable in terms of old ones: ontology extension can be substantive and has to be for humans.

- It is often assumed that discovering causes is discovering correlations, or laws relating observed phenomena (Hume). I'll try to show that some causal understanding goes deeper and is based on understanding of interacting structures – not necessarily all visible (Kant).

Could evolution have produced this? (Chappell&Sloman, 2007, IJUC)

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

Pre-configured competences:

are genetically pre-determined, though they may be inactive till long after birth (e.g. sexual competences), and their growth may depend on standard, predictable, features of the environment, as well as on DNA.
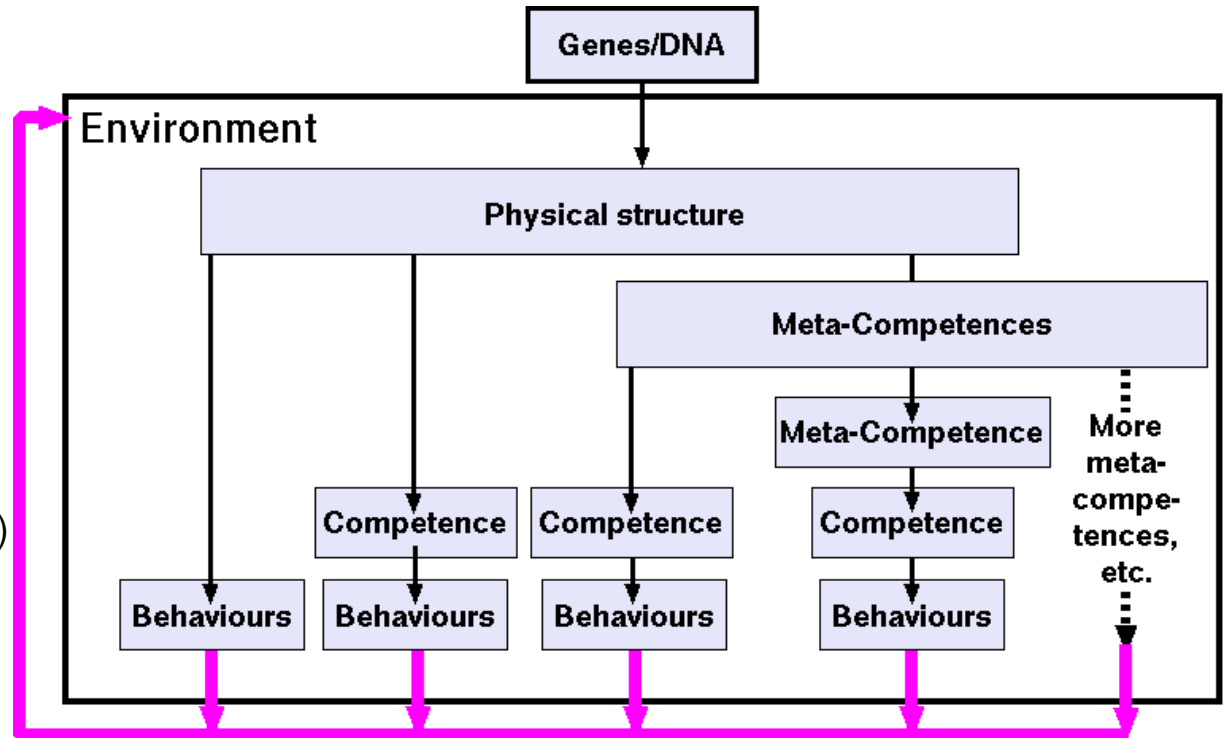
They occur towards the left.

Meta-configured competences:

(towards the right of the diagram) are produced through interaction of pre-configured or previously produced meta-configured competences with the environment (internal or external).

The environment changes the learning and development mechanisms.

Genes/DNA

Environment

Physical structure

Meta-Competences

Meta-Competence — More meta-compe-tences, etc.

Competence   Competence   Competence

Behaviours   Behaviours   Behaviours   Behaviours

(Chris Miall helped with the diagram.)

Evolution 'discovered' that speed of learning is increased by active intervention: it produced some species that discover many facts about the environment, and themselves, through creative exploration and play, in which ontologies, theories and strategies are developed, tested and debugged.

PERHAPS INFANTS THAT STARE LONGER AT SOMETHING ARE TRYING TO DEBUG A THEORY?

3. Two Kinds of Causal Understanding

This is the main topic I wished to discuss.

(Parts 1 and 2 provided some methodological and theoretical background.)

Analysing the concept 'cause' is one of the hardest problems in philosophy.

There are two main, very different, ways of viewing causation: Hume's and Kant's.

Usually philosophers argue about which one is correct.

I suggest that for children, scientists, some other animals and probably future robots also,

both concepts of causation are useful, in different ways, and at different stages in the development of understanding of some range of phenomena.

Conjecture, and Videos of human infants

Two ways of thinking about, representing and using causal connections.

- Humean causation is concerned with statistical relations between 'atomic' facts

- Kantian causation is concerned with deterministic relations between complex structures and complex processes involving those structures.

When multi-strand relationships change multi-strand processes occur.

Some processes are intelligible and predictable.

Others are merely instances of previously observed types of process.

Conjecture:

**The evolution of mechanisms for perception of 3-D processes,
such as grasping, twisting, levering, sliding, breaking, fighting
(and many more) led to the ability to understand
Kantian causation and the ability to do Kantian causal reasoning.**

Some infant/toddler videos are available here, discussed below
`http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/vid/`

NOTE: Felix Warneken's presentation suggested that young (pre-linguistic) toddlers, and also chimpanzees are able to do some Kantian causal reasoning about what underlies visible human behaviour.

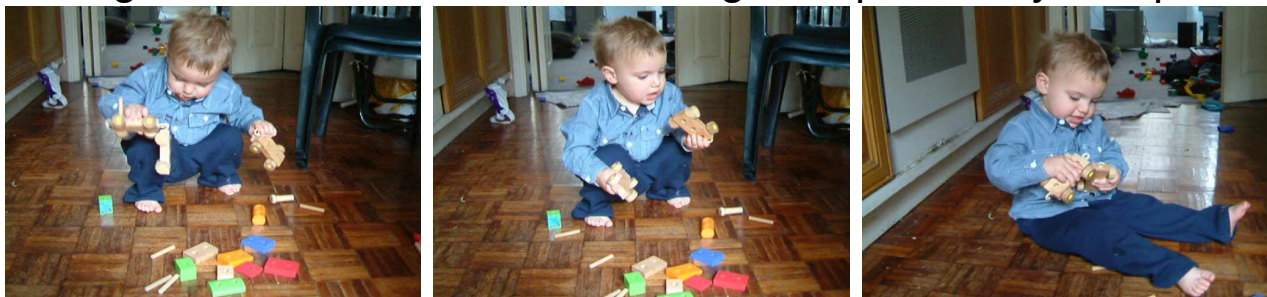# Causal competences: Yogurt, broom and train

Playing with yogurt and with ideas (11 months)

Pushing a broom – using competence and luck (15 months)
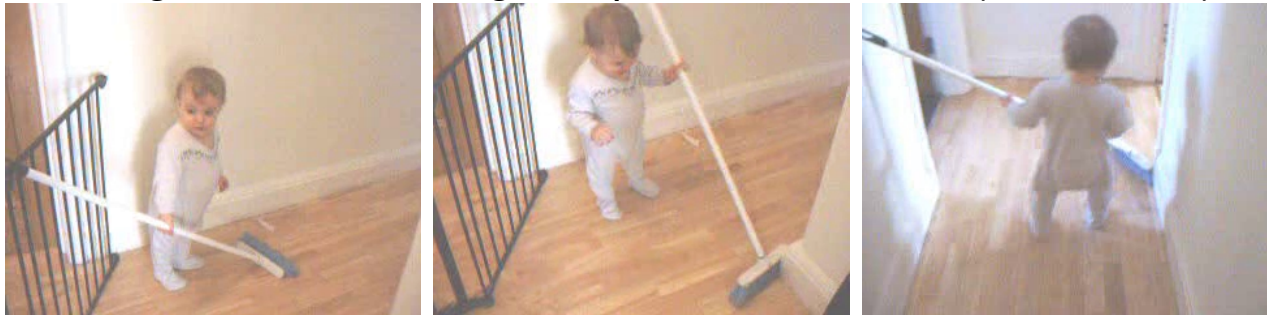
Failing to understand hooks and rings despite many competences (18 months)

The videos are here, along with several more.

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

The Child as Scientist: 1

Yogurt can be food for both mind and body in an 11 month baby.

Video available at

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

Hypothesis

Alongside the innate physical sucking reflex for obtaining milk to be digested, decomposed and used all over the body for growth, repair, and energy, there is a genetically determined information-sucking reflex, which seeks out, sucks in, and decomposes information, which is later recombined in many ways, growing the information-processing architecture and many diverse recombinable competences.

HOW ???

There is more discussion of this video in

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0603`

It is suggested that the child attempts to put yogurt onto the carpet and onto his thigh and fails because he lacks the ontology required to represent the role of the bowl of the spoon in transferring yogurt and does not realise that a rotation is required move the bowl away from being between the yogurt and the target location.

Later the child extends his ontology. How?

(He probably does not know he is doing any of the above.)

Child as scientist 3: Developing dynamical control systems and cognitive mechanisms in parallel

Pushing a broom – using competence and luck (15 months)

Video available at

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

Growing dynamical control systems

Some of what's going on is developing balance while walking and also developing control mechanisms for pushing the broom so that it moves as intended: both require continuous feedback in complex dynamical control systems.

Growing 3-D perception and planning mechanisms, using causal understanding

There is also perception of structural relations, e.g.
- back of broom handle caught between two rails – restricting sideways movement,
- broom obstructed by wall – restricting forward movement,
- wall at end of corridor – restricting forward movement,
- doorway appearing in right wall – allowing movement to right

Perceiving these things leads to goal-directed alteration of applied forces, e.g. to free the broom handle.

In some cases it can lead to formation of a future goal – e.g. to go through the doorway.

Child as scientist 3: Failing to deal with hooks at 19 months

1: Lifting two trucks makes the third disengage.

2-3: He picks it up with his left hand & shakes off the hanging truck with his right.

4: He notices the blank end & puts the truck down, rotating it.

5: He makes a complex backward move from crouching to sitting – while leaning forward to pick up the rotated truck.

6: He sees two rings.

7-9: He tries to join the rings, ignoring the hook, fails and gets frustrated, bashing trucks together and making an angry sound.

See the video   http://www.jonathans.me.uk/josh/movies/josh34_0096.mpg

Within a few weeks, he had learnt to see and use the hook-affordances. How? (Nobody saw how.)

The Child as Scientist: 4

- The idea that an infant, or possibly an older child, is like a tiny scientist investigating the world is often reinvented.

- It is obviously false if taken literally, for instance, because there are many conceptual, representational and mathematical tools used by scientists that are not available to a child, not even highly talkative and competent four-year-olds.

- A currently popular view, exemplified by work of Alison Gopnik and colleagues
    online here http://ihd.berkeley.edu/gopnik.htm

  is that young children (or at least their brains!) have the prerequisites for making causal inferences consistent with causal Bayes-net learning algorithms, which deal with conditional probabilities.

- On that view the concept of cause is viewed as concerned with correlations – Humean causation with conditional probabilities replacing universal correlations.

- Another view, implicit in Kant's critique of Hume, points to a deterministic, notion of causation concerned with structures and their interactions. On this view understanding causation is, at least in some cases, akin to proving, or at least understanding, mathematical theorems as in geometry.

- We suggest that the probabilistic/correlational (Humean) kind of causality is what most animals have, but humans and maybe a few others also have something deeper: a Kantian, deterministic, structure-based understanding of causality – the sort that drives deep science. To be a Kantian scientist is to be, in part, a mathematician.

Humean causation today: Bayesian nets

In recent years, developments of Humean ideas about causation have led many researchers to assume that causal information can be represented in the form of Bayesian nets, which sum up a collection of conditional probabilities – with or without numbers indicating the precise probabilities.

The diagrams can be interpreted as expressing Humean associations:

1. could represent: As cause Bs (e.g. failed brakes cause crashes)

2. could represent: As cause Bs and Bs cause Cs
(failed brakes cause crashes and crashes cause injuries)

3. alters that to: Bs can be caused either by A1s or by A2s or both

4. changes the above to represent A2s being direct co-causes of Cs, along with Bs

5. Adds Bs also cause Ds

Numbers can be added to the diagrams to express prior probabilities and conditional probabilities.

Experimental tests e.g. clamping some thing to prevent it from changing, or directly altering something to discover what does and does not change, can be used to evaluate a particular Bayesian net as a theory about a set of causal relationships.

Each link can be regarded as representing a true counterfactual conditional statement, or a conditional probability, if numbers are added.

(For a gentle introduction see: Steven Sloman, *Causal Models*, OUP, 2005)
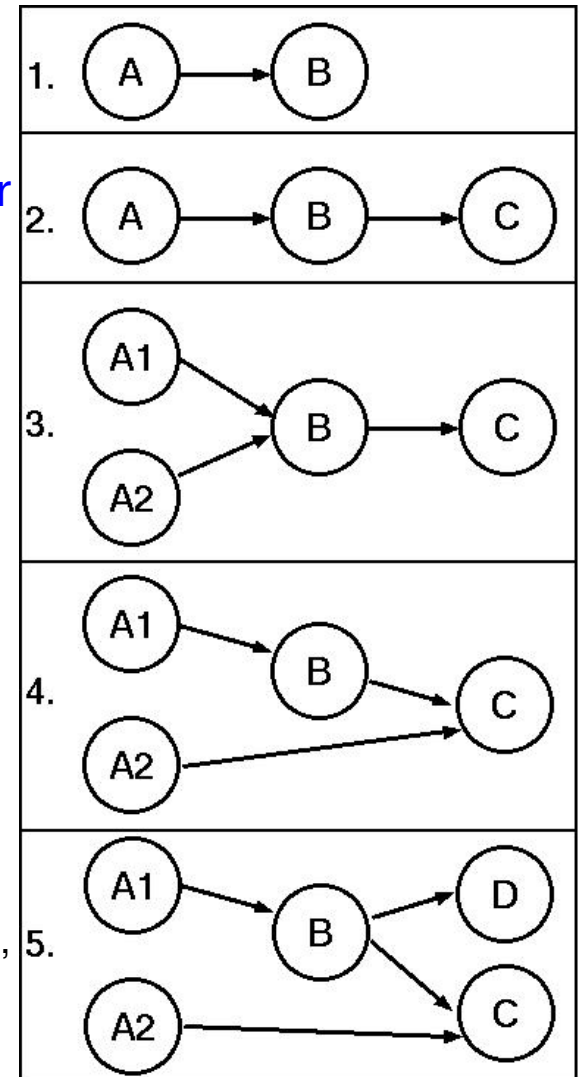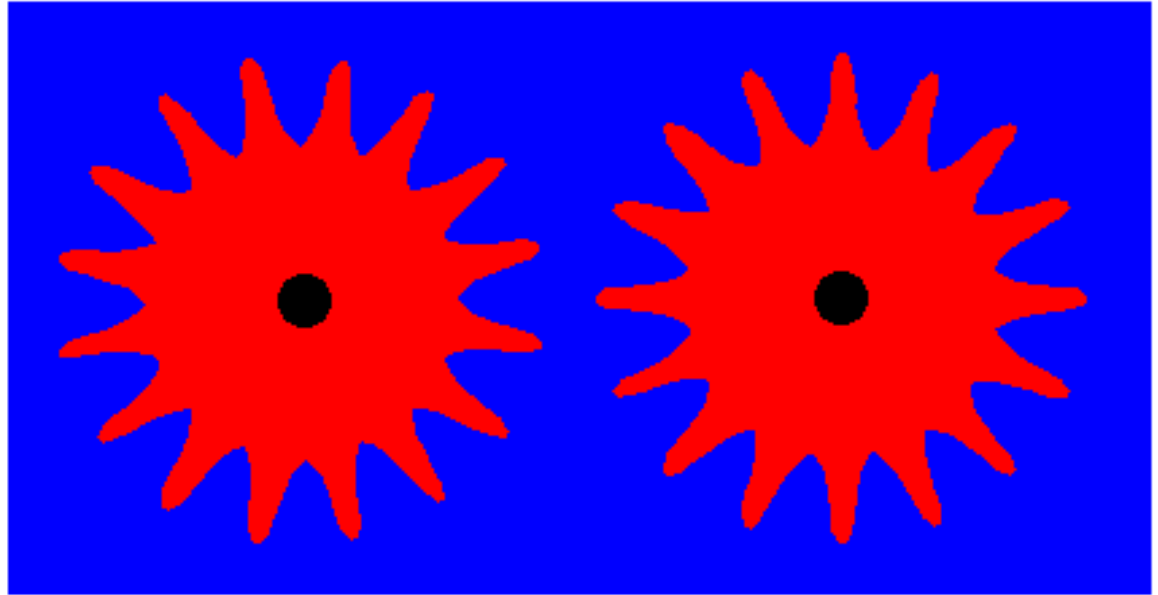
Causation in gear wheels (Humean)

Two gear wheels attached to a box with hidden contents.

Can you tell what will happen to one wheel if you rotate the other about its central axis?
Not on the basis of what you see.
You (or a child) can build a Bayesian net by experimenting.

- By experimenting you may or may not discover a correlation between the rotations – 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 that which way the second wheel rotates depends on the position of the knob or lever.
  (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, etc. etc.

Causation in gear wheels (Kantian)

Gear wheels closer together and meshed:

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.

What you can see includes the following:

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 both are rigid and impenetrable, then if the first tooth continues moving, it will push the other in the same direction, causing its wheel to rotate in the opposite direction.

The above is an example of a multi-strand process: many relationships change.

Compare understanding causation in a lever.

(We are not claiming that children need to reason **verbally** like this, consciously or unconsciously: explaining the forms of representation used is research that remains to be done.)
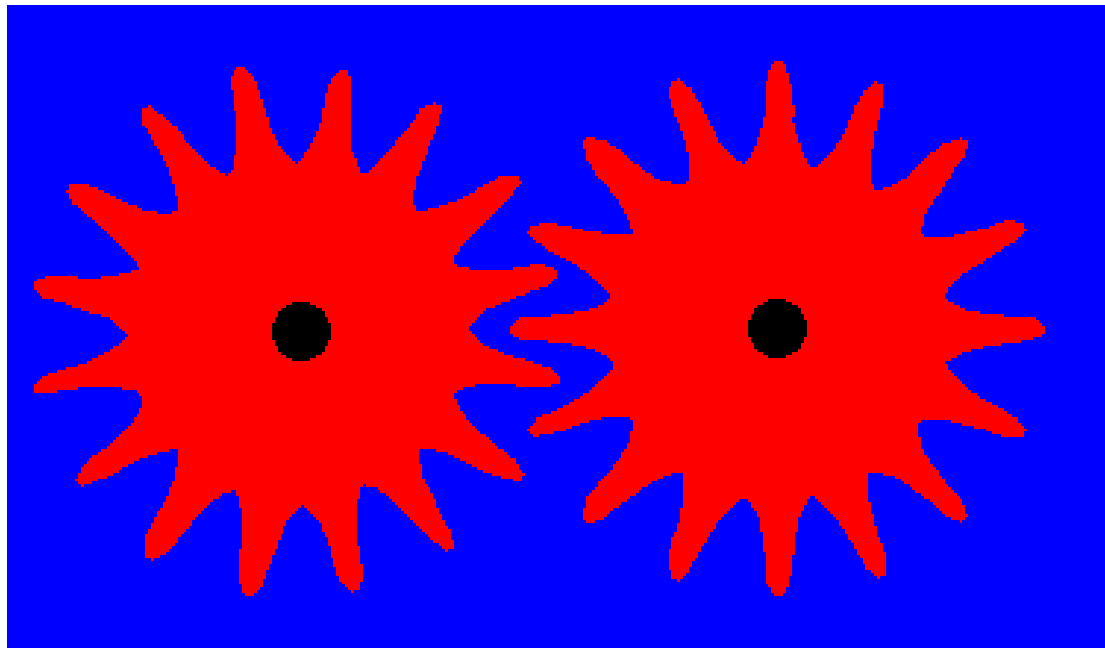
Understanding, simulating, proving

Kantian reasoning about effects is partly like running a simulation, except that it does not require precise details of the process to be simulated.

The simulation happens at a level of abstraction that is comparable with reasoning about a theorem in Euclidean geometry.

E.g. such and such a construction will produce a line that divides the triangle into two parts with the same area.

We don't need perfect precision in our diagrams or our simulations – because they are representations of processes, not replicas or models.

See Sloman, (IJCAI 1971), Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence

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

Not just geometry

Besides perceived shape, and interactions between shapes, the simulations that you run can also make use of unperceived constraints:

e.g. rigidity and impenetrability and the fixed location of the axle.

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

Constraints and processes using them need not be conscious, or expressed in linguistic or logical form.

There is no suggestion that a child can do all this from birth.

How this competence develops, and what information processing is involved, remains to be explained.

Use of Kantian causation is possible with partial knowledge

A theory using Kantian causation can be used to explain observations.

Suppose part of the configuration is covered, and it is observed that when the teeth visible on the left are moved down, the teeth on the right also move down, and reverse their motion when the teeth on the left are moved up.
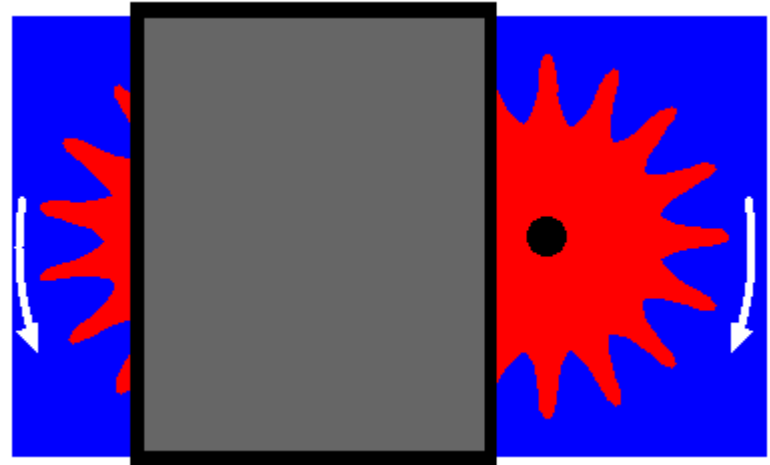
If you have a theory about kinds of stuff out of which objects can be made that are rigid and impenetrable, then you can combine that with a theory about shapes in a Euclidean geometry and possible motions in which parts of shapes move while others are fixed, like a wheel pivoted at its centre.

That theory of moving, rigid, impenetrable shapes has very many models, one of which could be a pair of meshed gear wheels each centrally pivoted.

Building and running a hypothetical model of the theory, including portions not observed, allows the observed phenomena to be explained, by being predicted, using the hypothesised explanatory facts and geometrical reasoning.

In this case, using the constraints of the theory you can work out that if the left wheel rotates clockwise the right wheel must rotate anti-clockwise – and *vice versa*. (No robot can do such things yet.)

NOTE: Theoretical concepts, like Rigid and Impenetrable used in the theory are NOT definable in their full generality in terms of what you can sense, but their roles in the theory, and their usability in making predictions give them meaning, for a user of the theory able to manipulate forms of representation corresponding to possible models. `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models`

The role of abduction

Finding causal explanations using Kantian causation is a form of abduction where the representation used is sometimes not logical but geometric.

NOTE:

Abduction is often incorrectly presented as a third form of reasoning alongside deduction (working out what must be the case, given some premises) and induction (working out what is likely to be the case, given some premises that do not suffice to support deductive inferences).

The key notion of abduction is not a process of drawing some conclusion by reasoning, but a process of formulating a conjecture which can be combined with other assumptions in a process of reasoning, e.g. to explain or predict something.

The construction of explanatory theories in science and everyday life requires abduction.

Abduction involves selecting candidates from a search space, e.g. the space of possible hypotheses (or geometric models) that could usefully and consistently be added to what is already known.

In general that is a huge space, and if the search is not constrained in some way it can be impossibly difficult.

In practice, the search for explanations of familiar phenomena in familiar contexts can be heavily constrained by what has already been learnt about those contexts and by what is already known – like finding an explanation for the motion of a gear wheel in the previous example.

The roles of abduction in philosophy and science are compared and contrasted here:

`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/logical-geography.html`

Some relevant papers on multimodal abduction by Lorenzo Magnani are mentioned at the end.

This is all about requirements – no designs yet

The examples presented so far merely give indications of what needs to be explained, in biology and psychology, and what needs to be implemented in future robots (e.g. future general-purpose domestic robots of the sort that inspired requirements analysis in CoSy).

At present nobody knows how to build working systems that have these capabilities and nobody knows how to explain how human and other animal brains are able to do Kantian causal reasoning, or even how 3-D structures and processes are represented in perception.

So I am saying neither that psychologists have some knowledge from which roboticists can benefit, nor that roboticists have some knowledge from which psychologists can benefit (though both are true), but that both need to recognize that they share major gaps in areas of common interest (e.g. understanding causation), and perhaps by working together they can both make faster progress.

Psychologists can document the variety of phenomena, especially intermediate cases in the development of visual, geometric, and causal competences, and associated learnt ontologies, and AI researchers (including neuroscientists using the design stance) can explore new forms of representation, algorithms and architectures inspired by the task of modelling all those cases.

Perhaps by working together we can understand and model the development of both Kantian and Humean causal understanding, and how each leads to the other.

NOTE:

Kantian causal reasoning plays a role in many aspects of human life, including humour, for instance in droodles, as illustrated in the next slide.

A Droodle: What do you see?

In many cases what you see is driven by the sensory data interacting with vast amounts of information about sorts of things that can exist in the world.

But droodles demonstrate that in some cases where sensory data do not suffice, a verbal hint can cause almost instantaneous reconstruction of the percept, using contents from an appropriate ontology.

Often understanding a droodle requires the ability to do Kantian causal reasoning about unseen mechanisms and processes: you can visualise a process producing the depicted result.

See also
http://www.droodles.com/archive.html

Verbal hint for the figure: 'Early worm catches the bird' or 'Early bird catches very strong worm'

Humean causation lacks the generative power to cope with cases like this.
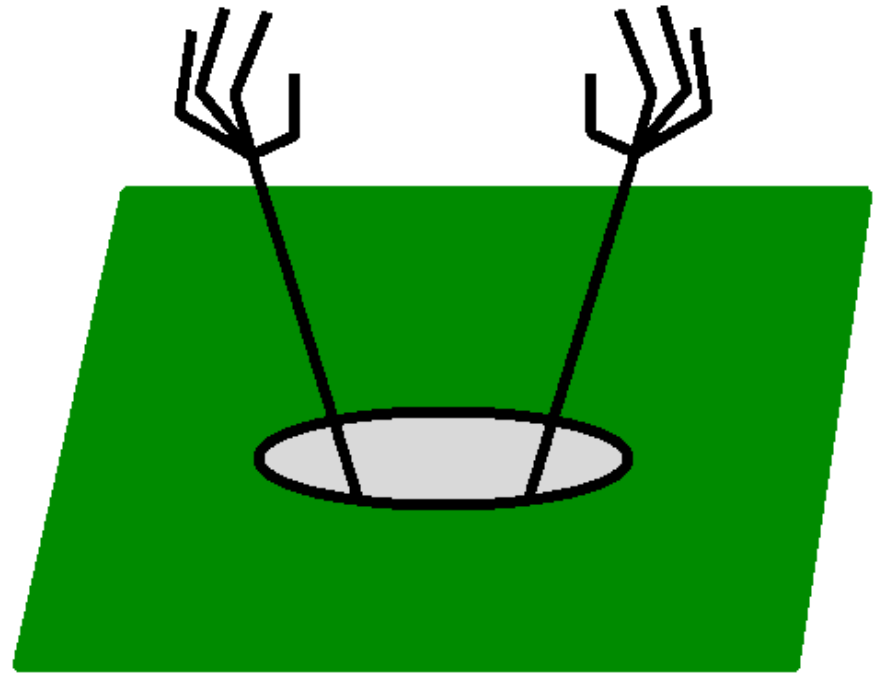
Many other jokes use Kantian causal reasoning

Many jokes and cartoons depend on the ability to perceive a situation and either to reason causally into the future to some calamity that is about to happen, or, less often to reason causally backward to some previous state.

This is a form of abduction: searching for an explanation.

Here's an example of the latter kind, from a book entitled *French Cartoons* edited by William Cole and Douglas McKee, published in 1955 by Panther Books.
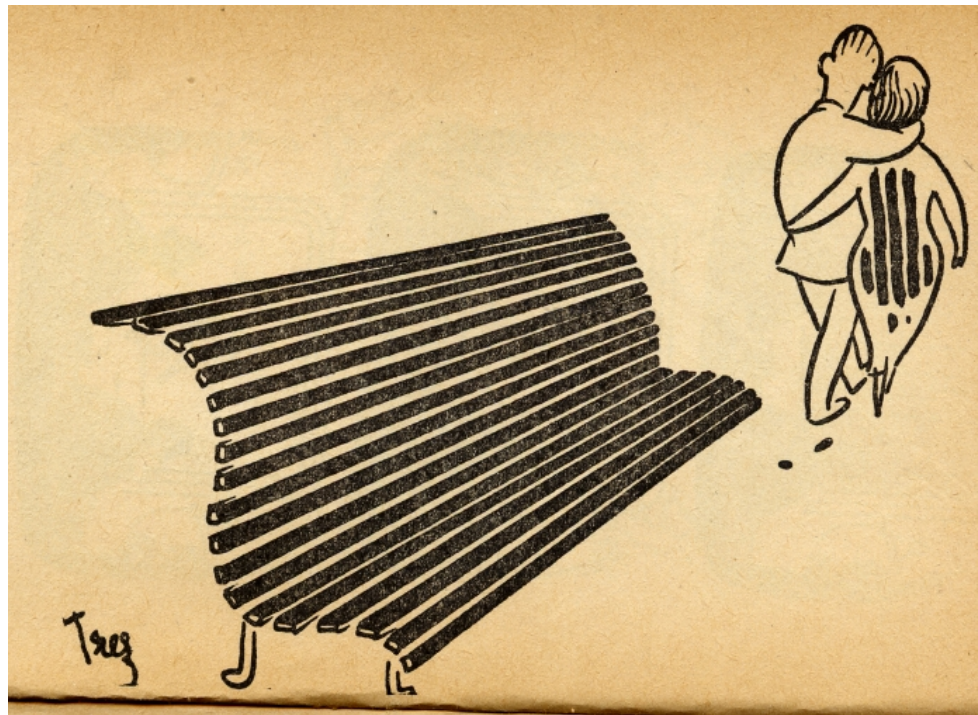
Limitations of Humean causation

Humean conceptions of causation, including the Bayesian net model, treats causation as a relation between events or states of affairs without representing the kind of detail required for Kantian causal understanding.

This does not use the ability of humans and some other animals to understand causation in terms of the operation of structured mechanisms, like gears, levers, strings, pulleys, where causation involves concurrent interaction between multiple parts.

Something other than Humean causation is needed to explain causal understanding based on our ability to see structured 3-D processes in which

- there are multiple relationships between objects and parts of objects

- different relationships change concurrently
  - some continuously
    (e.g. direction, distance, angular or linear velocity, acceleration, pressure, location of contact, etc.)
  - some discretely
    (e.g. contact beginning or ending, containment changing, collision impending or not, something moving into or out of view, occlusion, containment, obstruction starting or ending, etc.)

- some changes are necessarily connected with others
  (e.g. because of the geometry of the situation: if X continues in the same direction it must hit Y)

- There are additional, non-geometric constraints that come from physical properties of the materials involved, e.g. rigidity, impenetrability, elasticity, density, stickiness, etc.

The over-rated "DO" operator: for surgery on nets

There is a long philosophical tradition emphasising the role of experimentation in discovering what causes what – and this can lead to a spurious distinction between learning from observation and learning from doing.

This idea is sometimes enshrined in the notion of a DO operator on a Bayesian net.

E.g. If A-events and B-events often occur together, and never alone, and a C-event follows immediately, you may conjecture that A and B together cause C-events.

If you can "DO" A while you "DO" preventing B from occurring, and you then observe that C occurs, you have learnt that your causal net can have a link from A to C without B being involved: the net needs "surgery" (Pearl).

But there is no logical difference between noticing this when you have done the experiment and noticing it when someone else does the experiment or noticing it when the experiment is produced without any human agency by the wind blowing in such a way as to produce A without producing B.

- The main difference between passive observation and active intervention is speed of learning: you may have to wait a long time to observe something passively.

- Active intervention can speed up discovery of new relations and testing of hypotheses but it does not provide stronger evidence than passive observation: it merely provides more, faster.

- Of course people may treat the results of their own interventions as somehow providing better evidence, but that is just one of many ways in which human reasoning can be fallacious.

- The study of how people actually reason (psychology) should not be confused with the study of valid reasoning (logic and mathematics).

Advantages of Kantian causation (when available)

- If the only way you can find out what the consequence of an action will be is by trying it out to see what happens, you can acquire knowledge of causation based only on observed correlations.

  This is 'Humean causation' – constant conjunction, enriched by including conditional probability distributions and Bayesian nets.

- Sometimes you don't need to find out by trying because you can work out consequences of the structural relations (e.g. by running a simulation that has appropriate constraints built into it): You are then using Kantian causation, which is deterministic, structure-based and generative:

  It supports understanding of novel situations, and designing new actions and machines.

- As children learn to understand more and more of the world well enough to run deterministic simulations they learn more and more of the Kantian causal structure of the environment, including how changing or moving spatial structures interact.

- Typically in science causation starts off being Humean until we acquire a deep (mathematical) theory of what is going on: then we use a Kantian concept of causation.

- Not all Kantian causal reasoning uses geometry: Science and mathematics often produce new formalisms for reasoning with.
  Newton and Leibniz invented new ways of reasoning about interactions between position, rate of change of position, rate of change of rate of change, etc., using the new formalism of calculus.

Kant's most famous example: viewing a house

What will you see when you move round a house, in one direction or the other; or when you move from top to bottom of a house, or from bottom to top? See Kant, The Critique of Pure Reason Second Analogy

The example makes the point that if you understand some of the geometric and topological structure of the environment you can deterministically predict consequences of various movements.

- Actions that alter your location cause different parts of the house to change their appearance or become visible or invisible in systematic ways, which depend on the structure of the house.

- Some changes are continuous (e.g. continuous foreshortening) and some discrete (e.g. as a previously invisible face of the house becomes visible, or vice versa): and both sorts can occur at the same time.

- As with causation in levers, pulleys, or engaged gears, this is a collection of effects that are inferrable on the basis of an understanding of the structure of the situation: you can work out the effects of different combinations of movements.

    As illustrated in Sloman (1971) 2nd IJCAI (Reference above)

- This is closely related to our ability to do mathematical reasoning, especially geometrical and topological reasoning.

Compare the use of 'aspect graphs' in computer vision, and Minksy's 'Frame systems' (1973).

For more detailed discussion of Kant's views see, for example,

http://www.trinity.edu/cbrown/modern/litrev/Kant-causation.html
Kant's Second Analogy: Literature Review by Melanie Butler (?)
on Curtis Brown's web site: http://www.trinity.edu/cbrown/

Can non-human animals use Kantian causation?

This was one of the topics discussed at WONAC 2007

International Workshop on Natural and Artificial Cognition
Oxford, June 2007 (Funded by NSF and euCognition Network)

The aims and contents overlapped with the CoSy MofM workshop but with emphasis more on non-human animal cognition.

Information about the workshop is available here.

```
http://tecolote.isi.edu/~wkerr/wonac/
http://tecolote.isi.edu/~wkerr/wonac/program.html
```

Jackie Chappell's talk at WONAC discussed some of the problems of doing empirical research on understanding of causation in animals.

See her slides – two versions:

```
http://www.cs.bham.ac.uk/~jmc/chappell-wonac-slides-screen.pdf
```
(For screen viewing (with hyperlinks))

```
http://www.cs.bham.ac.uk/~jmc/chappell-wonac-slides-print.pdf
```
(for printing)

Post workshop versions of our slides, and further discussion of various levels of sophistication in Humean and Kantian causal competence can be found here, along with some more references:

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

Summary: simple forms of Humean (correlation-based) causal competence abound in animals.

Evidence of Kantian causal competence is rarer and more controversial, but some primates and birds (e.g. corvids) seem to have some level of Kantian understanding in at least some contexts.
  E.g. It may be required for some kinds of nest construction.

Similar questions can be asked about human infants and toddlers, and can be equally hard to answer.

Tamarins at work

At the WONAC Workshop, Jackie Chappell talked about the following

Hauser, M.D., Kralik, J. & Botto-Mahan, C. (1999).
Problem solving and functional design features: Experiments with cotton-top tamarins (Saguinus oedipus). *Animal Behaviour* 57: 565-582.

Available online
`http://www.wjh.harvard.edu/~mnkylab/publications/learnconcepts/problemsolving.pdf`

They conclude

...First, cotton-top tamarins readily solve means-end tasks. Second, in solving such tasks, the tamarins attend to the functionally relevant features. Featural transformations that do not affect an object's functionality were readily tolerated with regard to performance in the means-end task. This shows that tamarins can discriminate between objects that show signs of good design and those that do not, and can use this knowledge to select an appropriate object for solving a problem. Third, some featural transformations that were clearly relevant to the functionality of the task proved difficult for the tamarins and required explicit training. Thus, although the tamarins appear to have a greater capacity for creating learning sets than revealed by some early experiments on the closely related marmoset (Miles & Meyer 1956; see also Rumbaugh et al. 1996), there are limitations. At present, it is not clear whether these limitations arise due to their relatively poor dexterity, to problems in making fine-grained perceptual discriminations (e.g. narrow gaps in the connected problem), to conceptually mediated problems associated with understanding means-end problems (e.g. understanding abstract-relational concepts), or to some combination of these factors. .... Fourth, although the tamarins were able to inhibit some actions, some problems proved difficult because of their inability to inhibit a reaching response. Future work must assess the degree to which problems of inhibition over-ride a species' capacity to solve problems conceptually and the extent to which interspecific differences in the prefrontal cortex contribute to interspecific differences in problem-solving ability.

The summary assumes that the tamarins have a fairly rich ontology.

We can ask whether that ontology was pre-configured or meta-configured (see slide 20).
Whatever the answer, further questions arise, e.g. about precisely what information is represented, how the information is represented, how the representations are manipulated, how the results are used, etc.

These questions would have to be answered in designing robots with similar capabilities.

Don't ask which animals or machines understand causation

There is no fixed, sharp, distinction between understanding and not understanding causation.

Instead we can distinguish various of types of competence related to each of Humean and Kantian causation, as shown here

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/causal-competences.pdf`

and then ask:

- which animals have which subset,

- at what age or in what order they typically develop,

- how they evolved,

- how those competences are acquired or extended by individuals,

- which are necessary precursors for others,

- what forms of representation, mechanisms and architectures support them

- what the trade-offs are between alternative sets of competences and alternative implementations,

That way we can replace futile debates, about how to label phenomena, with productive research, on what sorts of competences different animals have, what the implications of those competences are, and what mechanisms can explain them.

We can do that for many debates about cognition and development:

Namely, replace ill-defined and poorly motivated dichotomies with analysis of spaces of possible designs and corresponding niches to be analysed, compared, explained, modelled, ...

No robots have Kantian causal understanding yet

Current work on robots and AI systems has not addressed the ability

to acquire and use Kantian understanding of spatial structures and processes and their interactions,

or to learn about different kinds of "stuff" and their properties

e.g. rigidity, impenetrability, elasticity, hardness, softness, etc.

That's mostly because

(a) those working on vision have focused largely on recognition, rather than understanding structure,

(b) researchers have not noticed the need to give machines Kantian causal understanding,

though some early work on "qualitative reasoning", e.g. by Kuipers, moved in this direction. See
`http://www.cs.utexas.edu/~qr/papers-QR.html`

(c) researchers have not given machines the ability to think and reason about their causal reasoning

(d) giving machines an understanding of 3-D spatial interactions is very difficult (except in very special cases).

AI work on fault diagnosis in circuits could be regarded as an exception to (b), and some early work on "naive physics" (Pat Hayes) and "qualitative physics" long ago.
(But see the ongoing work by Ken Forbus and colleagues and work on qualitative reasoning by Tony Cohn.)

Some limitations of current AI

Alas, current AI systems, including robots, have only a small set of the previous types of causal competence:

e.g. they may be able to do things, but, when doing something, don't know

what they do, why they do it, how they do it,
what they did not do but could have done
why they did not do it differently,
what would have happened if they had done it differently,
etc.

They also cannot tell whether someone else is doing it the same way or a different way,
and if something goes wrong explain why it went wrong.

All this is related to lacking Kantian causal competences (including lacking exosomatic ontologies and representations, as explained below).

Learning to do causal reasoning in different domains

The ability to work out consequences requires learning to build manipulable representations (information bearers) with appropriate structures, appropriate permitted changes, and appropriate constraints.

In some cases (e.g. geometric reasoning) the manipulations can be construed as more or less abstract simulations of changes whose effects are being predicted.

Compare: Kenneth Craik, *The Nature of Explanation,* Cambridge University Press, 1943,

What structures, changes and constraints are appropriate depends on what is being reasoned about: geometric reasoning about the rotation of a rigid gear wheel (e.g. one made of steel) is not the same as geometric reasoning about the rotation of something soft and malleable, e.g. putty or plasticine, or a tornado.

Appropriate constraints should ensure the right counterfactual conditionals are true as the transformations are done. (It is not obvious how to implement this!)

Work still to be done: Analysis of detailed representational, algorithmic, mechanistic and architectural requirements to support such learning, and how the ontology involved grows.

Part of the point of the CoSy Robot project is to investigate these issues, especially the requirements for human-like competence, which we need to understand before we can build designs or implementations, though the process of designing and implementing can help the process of understanding requirements.

But the problems are very hard and progress is very slow.
Moreover we don't think neuroscientists have discovered appropriate mechanisms either (in 2007).

For more detail on a theory of vision as involving running of simulations see
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505`

We cannot do it all from birth

Causal reasoning adults find easy can be difficult for infants.
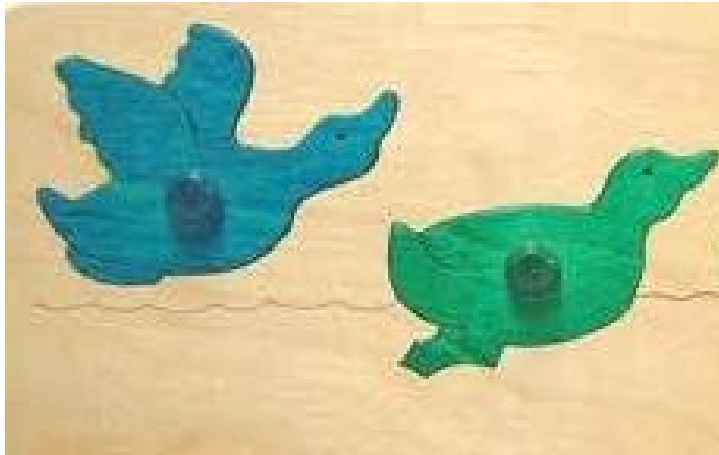
A child learns that it can lift a piece out of its recess, and generates a goal to put it back, either because it sees the task being done by others or because of an implicit assumption of reversibility. At first, even when the child has learnt which piece belongs in which recess there is no understanding of the need to line up the boundaries, so there is futile pressing. Later the child may succeed by chance, using nearly random movements, but the probability of success with random movements is very low. (Why?)

Memorising the position and orientation with great accuracy will allow toddlers to succeed: but there is no evidence that they have sufficiently precise memories or motor control.
Stacking cups compensate for that partly through use of symmetry, partly through sloping sides, so they are much easier.

Eventually a child understands that unless the boundaries are lined up the puzzle piece cannot be inserted. Likewise she learns how to place shaped cups so that one goes inside another or one stacks rigidly on another.

Conjecture:

Each such change requires the child to extend its ontology for representing objects, states and processes in the environment, and that ontology is used in a mental simulation capability. HOW?

Learning about kinds of causation

The process of extending causal competences is not continuous (like growing taller).

- The child has to learn about new kinds of
  - objects,
  - properties,
  - relations,
  - process structures,
  - constraints
  - forms of representation ,...

- and these are different for
  - rigid objects,
  - flexible objects,
  - stretchable objects,
  - liquids,
  - sand,
  - mud,
  - treacle,
  - plasticine,
  - pieces of string,
  - sheets of paper,
  - construction kit components in Lego, Meccano, Tinkertoy, electronic kits...

We don't know how many different things of this sort have to be learnt, but it is easy to come up with very many significantly different examples.

There is no fixed order in which things have to be learnt: there are many dependencies but not enough to generate a total ordering – each learner finds routes through several partially ordered graphs.
Think also of learners with physical and cognitive disabilities and abnormalities.

CONJECTURE

In the first five years

- a child learns to run at least hundreds, possibly thousands,
  of different sorts of simulations,

- using different ontologies
    with different materials, objects, properties, relationships, constraints, causal interactions —
        some opaque and Humean others transparent and Kantian.

- and throughout this learning, perceptual and representational capabilities are
  extended by adding new sub-systems to the visual architecture, including new
  simulation capabilities, and to other parts of the architecture

In the case of humans, things available to be learnt keep changing from one generation to another:
provision of new kinds of playthings based on scientific and technological advances is a major form of
communication across generations.

Some toddlers today do things that require ontologies undreamed of even by their grandparents.

Some more examples are available in

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601`

COSY-DP-0601 Orthogonal Competences Acquired by Altricial Species (Blanket, string and plywood).

Contrast that with adjusting parameters in a dynamical control system using hundreds or
thousands of practice trials.

Pushing and pulling

As toddlers learn to push, pull and pick things up, they find that some things 'hang together': if you move a part other parts move. But the growing ontology, and mechanisms for representing actions and their perceived effects need to allow for things that hang together in different ways.

If a group of bricks is lying on the floor, pushing a brick on the boundary towards the centre can make the whole group move, whereas pulling it in the opposite direction moves no other brick.

If you push the edge of a blanket towards the centre most of the blanket does not move, whereas if you pull the edge away from the centre the blanket follows (in an orderly or disorderly fashion, depending on how you pull, with one or two hands, etc.).

A sheet of paper the same size as the blanket will typically behave differently: pushing and pulling will move the whole sheet, but the effect of pushing will be different from pushing a pile of bricks (in what ways?) and the effect of pulling will be different from pulling the blanket (in what ways?).

What they have in common includes the fact that if a toy is resting on the blanket or sheet of paper, pulling the edge towards you brings the toy closer too, whereas if you pull too fast, or if the toy is on the floor near the far edge, pulling will not have that effect. Why not?

The child's ontology has to allow not only for different kinds of stuff (cloth, wood, paper, string, etc.), but also different ways in which larger wholes can be assembled from smaller parts: which requires a grasp of relations of different kinds, including 'multi-strand relations', and the 'multi-strand processes' that occur during changes in multi-strand relations, as discussed in

```
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0507
```

Some of the understanding of causation in such processes may start off Humean (i.e. using only conditional probabilities) and then as the ontology is enriched to include properties like *rigid, flexible, impenetrable, elastic, inextensible,* and these are combined with shape and spatial relations, the understanding can become more Kantian, i.e. structure-based, generative and deterministic, supporting more creative exploration and discovery.

Blanket and String

If a toy is beyond a blanket, but a string attached to the toy is close at hand, a very young child whose understanding of causation involving blanket-pulling is still Humean, may try pulling the blanket to get the toy.

At a later stage the child may either have extended the ontology used in its conditional probabilities, or learnt to simulate the process of moving X when X supports Y, and as a result does not try pulling the blanket to get the toy lying just beyond it, but uses the string.
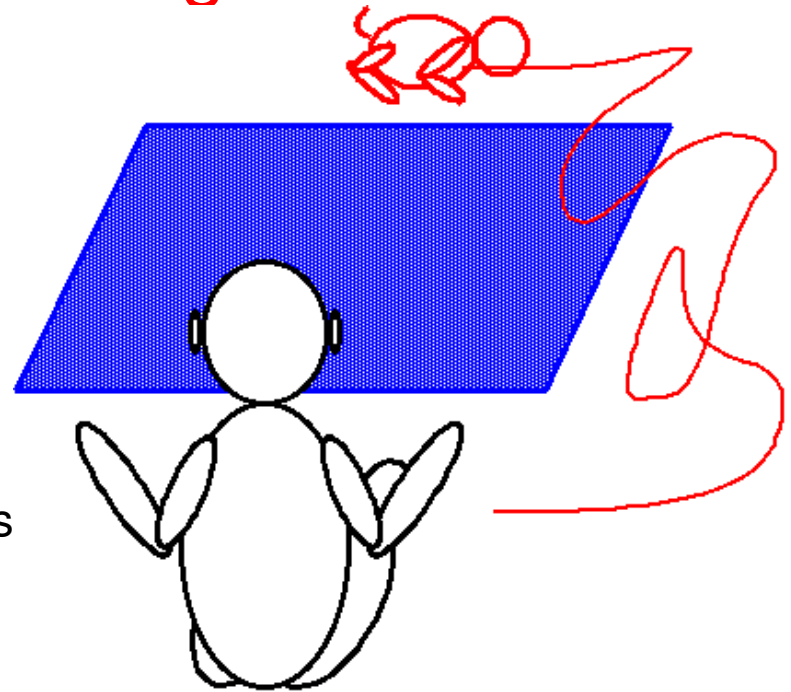
However the ontology of strings is a bag of worms, even before knots turn up.

Pulling the end of a string connected to the toy towards you will not move the toy if the string is too long: it will merely straighten part of the string.

The child needs to learn the requirement to produce a straight portion of string between the toy and the place where the string is grasped, so that the fact that string is inextensible can be used to move its far end by moving its near end (by pulling, though not by pushing).

Try analysing the different strategies that the child may learn to cope with a long string, and the perceptual, ontological and representational requirements for learning them.

Ontologies for getting at something

Understanding varieties of causation involved in learning how to get hold of a toy that is out of reach, resting on a blanket, or beyond it.

Some things to learn through play and exploration

Toy on short blanket Grab edge and pull

Toy on long blanket Repeatedly scrunch and pull

Toy on towel Like blanket

Toy on sheet of plywood
Pull if short(!!), otherwise crawl over or round it

Toy on sheet of paper Roll up?
(But not thin tissue paper!)

Toy on slab of concrete Crawl over or round

Toy at end of taut string Pull

Toy at end of string with slack Pull repeatedly

String round chair-leg Depends

Elastic string
?????

See this discussion of learning orthogonal recombinable competences
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601`

It takes a lot of learning to develop all the visual and reasoning competences required for seeing and understanding these affordances – including visualising what would have happened if you had done something different, or if someone else were to move something.
Our spatial and visual competence goes far beyond actually doing.
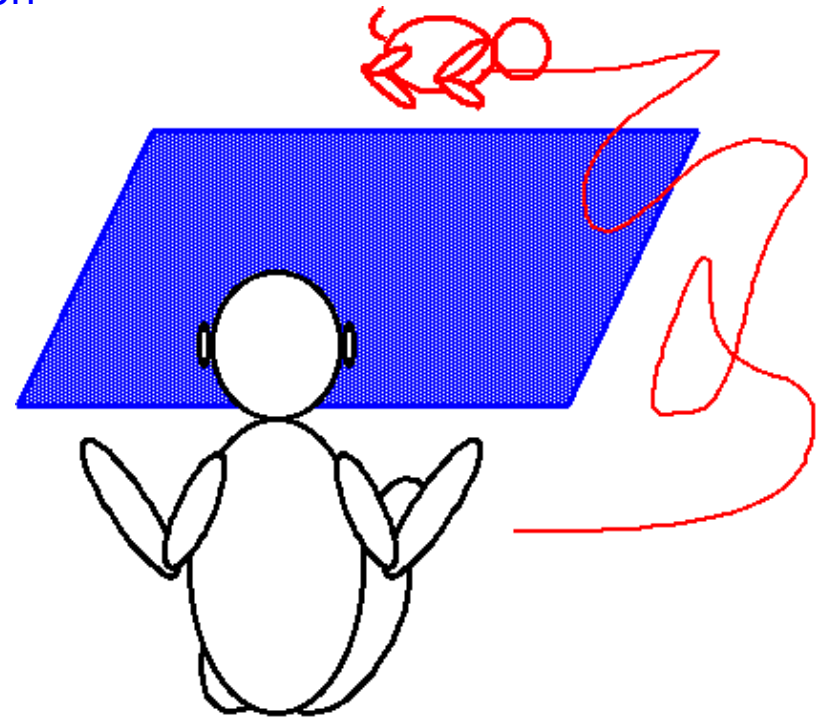
Cloth and Paper: Learning never ends

You have probably learnt many subtle things unconsciously, some as an infant or toddler, some later on, about the different sorts of materials you interact with (e.g. sheets of cloth, paper, cardboard, clingfilm, rubber, plywood).

That includes different ways in which actions can and cannot distort their shape.

Lifting a handkerchief by its corner produces very different results from lifting a sheet of printer paper by its corner – and even if the handkerchief had been ironed first (what a waste of time) it would not have behaved like paper.

Most people cannot simulate the precise behaviours of such materials mentally but we can impose constraints on our simulations that enable us to deduce consequences.

In some cases the differences between paper and cloth will not affect the answer to a question, e.g. in questions about results of folding processes that depend only on shape, not material.

Simulating motion of rigid, flexibly jointed, rods

A Kantian example: on the left, what happens if joints A and B move together as indicated by the arrows, while everything moves in the same plane? Will the other two joints move together, move apart, stay where they are. ???
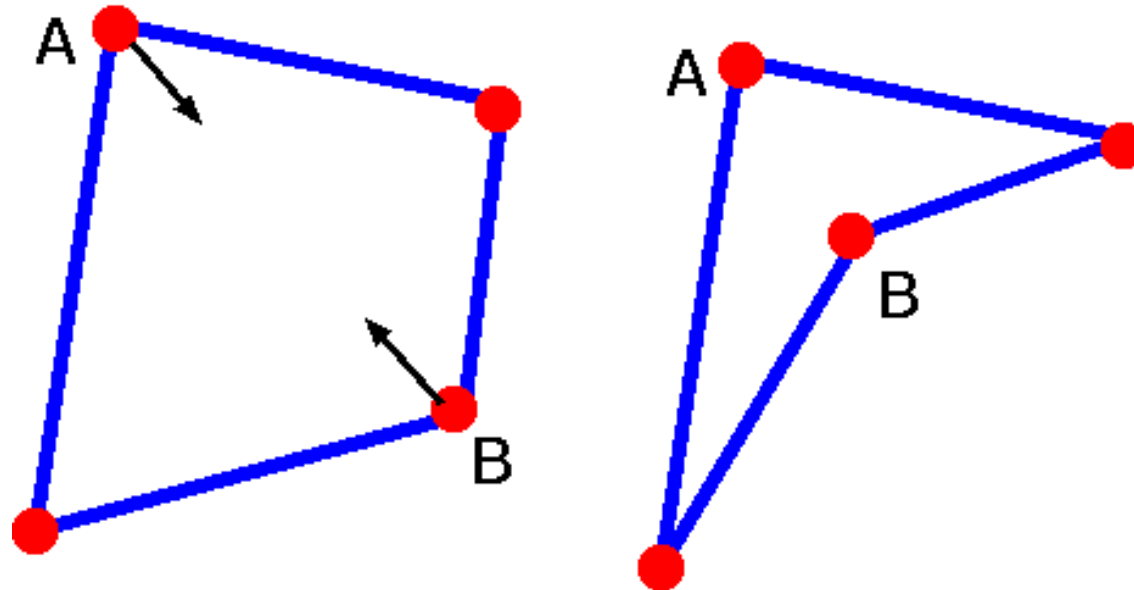
- What happens if one of the moved joints crosses the line joining the other two joints?
- This task is harder than the gears task (why?).
- We can change the constraints in our simulations: what can happen if the joints and rods are not constrained to remain in the original plane?
- What has to develop in a child before such tasks are doable?

Visual reasoning about something unseen

If you turn the plastic shampoo container upside down to get shampoo out, why is it often better to wait before you squeeze?

In causal reasoning we often use runnable models that go beyond the sensory information: sometimes part of what is simulated cannot be seen –

a Kantian causal learner will constantly seek such models, as opposed to Humean (statistical) causal learners, who merely seek correlations.

Note that the model used here assumes incompressibility and viscosity.

The model can also explain the fact that as more of the shampoo is used up you have to wait longer before squeezing.
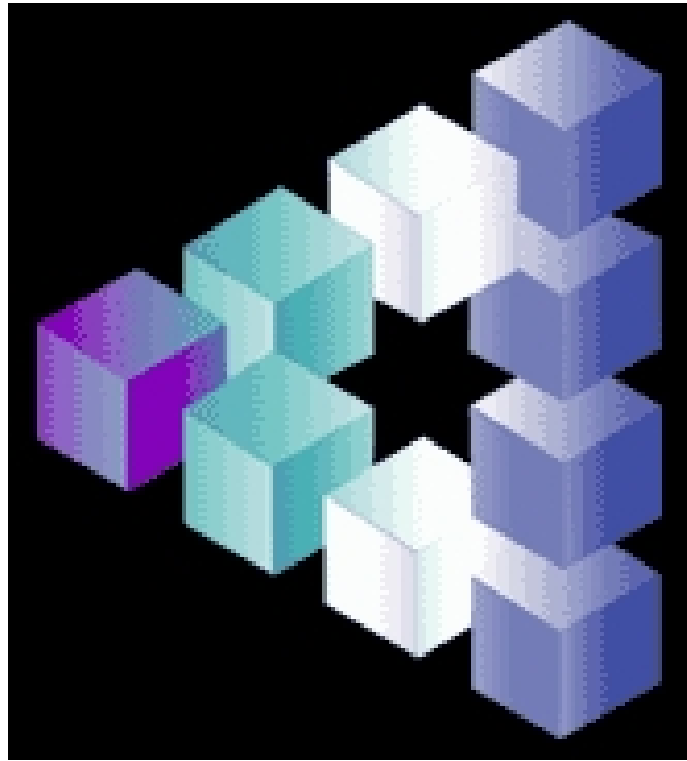
The next slide shows that the word 'model' as used here allows inconsistency.

So they cannot be models in the normal sense.
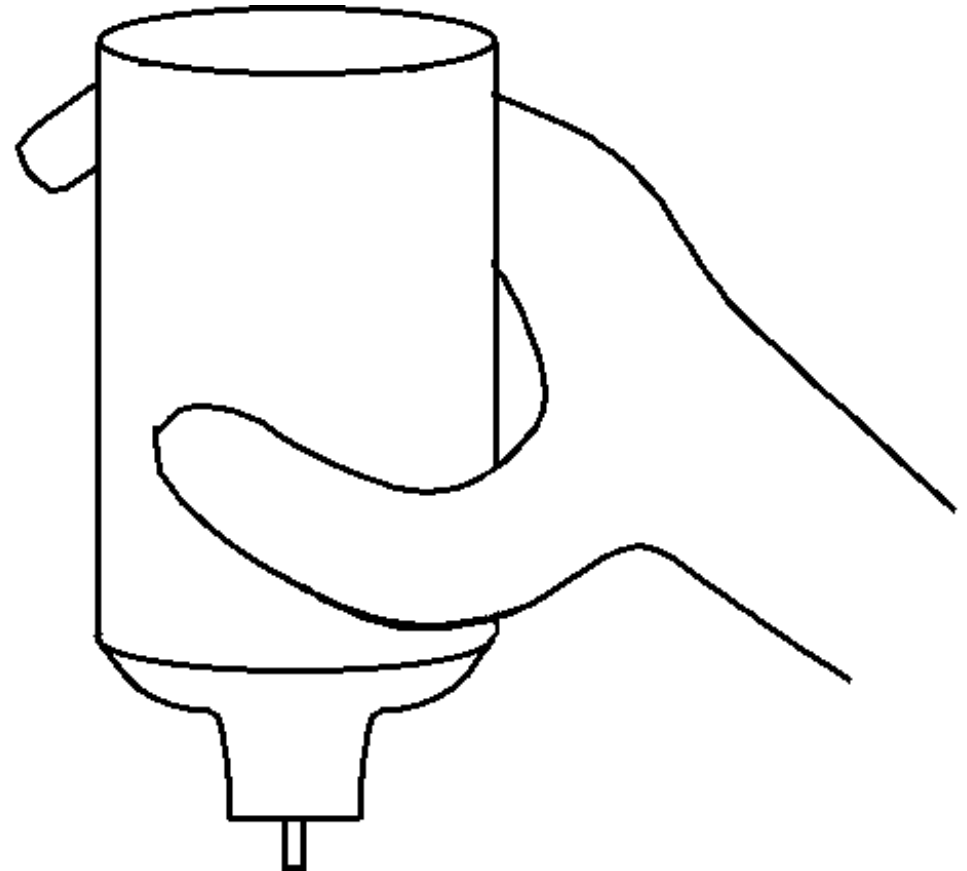
Note: a percept does not need to be consistent

A nice picture by Reutersvard – before Penrose

What are the implications of what you see?

Think of all the things you can do with or between the little cubes.

Seeing does not require a consistent percept

A model (in the normal sense) cannot be inconsistent: but a model of this scene would have to be.

Can you build what you see in a picture?

Children who learn to play with construction kits, e.g. meccano, lego, tinkertoy, Fischer-technic, plasticine, ... often learn to make things on the basis of 2-D pictures, which may be photographs of the actual models or sketches.

Normally, if a picture shows a configuration made of a small number of rectangular blocks and the construction kit consists of nothing but rectangular blocks, including the shapes shown in the picture, then it is easy to select a suitable collection of blocks from the kit and arrange them as shown.

The configuration shown is derived from a 1934 picture by Swedish artist Oscar Reutersvard.

> Someone trying to build the configuration depicted might keep trying, starting from different sub-configurations, and always fail, without understanding why. (Like the child mystified as to why pressing a picture into its recess does not make it go in?)

> Someone with deeper insight into structural relations will be able to see why it is impossible to do unless at least two blocks are omitted (concerning which there are alternative choices).
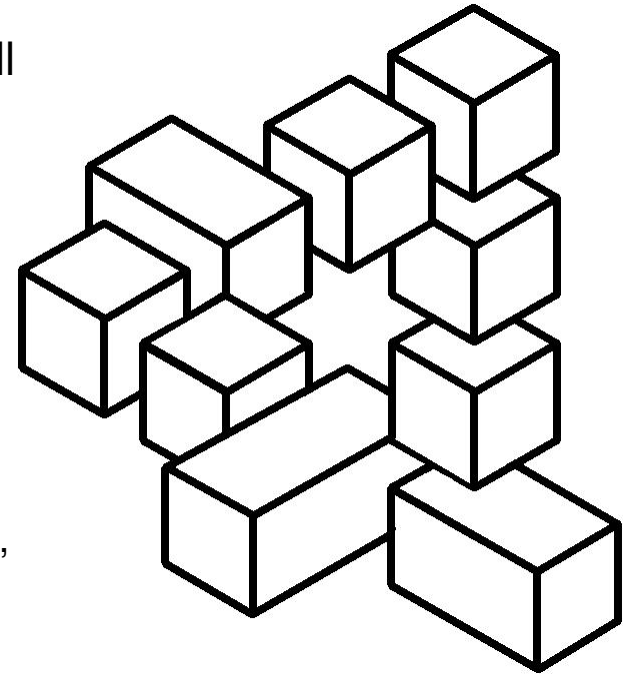
What has to change in a child's mind, and brain, to enable the impossibility to be detected and understood? In particular, what extensions are required to the forms of representation used, the ontology, the manipulative mechanisms?

Visual percepts cannot be models

Models cannot be inconsistent

However if percepts are made up of fragments combined in a manner that does not correspond to full spatial integration then inconsistencies are possible.

E.g.

- A is bigger than B
- B is bigger than C
- C is bigger than A

or, more plausibly, a large collection of proto-affordances of different sorts, spatially located.

Why might the use of such a fragmented, though spatially related, collection of distinct interpretations of portions of the scene be desirable?

Because the very same scene needs to be perceivable in different ways, depending on current goals, interests, etc.

So it must be possible to switch different items of information in and out of the percept.

E.g. different affordances, different relationships, low level or high level details.

This is one form of attention switching.

Sensorimotor knowledge and Kantian causation

The extreme 'sensorimotor' approach: everything an organism or robot needs to know is expressed in sensorimotor contingencies: conditional probabilities linking patterns in sensory and motor signals, at various levels of abstraction.

(Using a multi-modal "somatic" ontology).

It can also be presented as a theory of consciousness.
O'Regan, Noë, Hurley and others:
    http://lpp.psycho.univ-paris5.fr/tikiwiki/tiki-index.php

There are strong connections with 'symbol grounding' theory, criticised here:
    http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models

Sensorimotor models characterise how most simple organisms represent the environment, e.g. microbes and insects most of the time, and many others

But using representations of 3-D spatial structures, processes, relationships, and unobservable but inferrable properties of kinds of stuff (as required for the Kantian conception of causation), goes beyond sensorimotor representations and uses an a-modal exosomatic ontology.

Selection pressure may have favoured exosomatic ontologies and forms of representation because they drastically reduce the combinatorics of representing grasping and other actions performed by hands or claws that move independently of each other and of the eyes (unlike mouth and beak).

However, training of rapid responses can produce "compiled" fast-acting sensorimotor representations of special cases of more general competences, if those special cases occur often enough: So, the same generic competence may use different forms of representation in different contexts

This is true even of abstract symbolic skills, e.g. in mathematics, musical sight-reading, language comprehension and generation, as well as game playing, athletics, car driving, etc.

More on somatic/exosomatic ontologies

Here are some presentations and discussion papers on the limits of the use of sensorimotor, somatic, ontologies.

`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/nature-nurture-cube.html`
   How to see a rotating 3-D wireframe cube (Necker cube).

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601`
   Learning orthogonal recombinable competences.

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603`
   Critique of sensorimotor contingency theory.

Could the portia spider perform her amazing feats if she had only representations of her sensory and motor signals and relations between them?

See
   `http://www.freerepublic.com/focus/f-chat/1640513/posts`

Conjecture:

   Much of the fuss about so-called "mirror neurons" is really about evolution's solution to the problem of defeating the combinatorial complexity of somatic ontologies by representing some things in the environment independently of the perceiver's sensory and motor signals.

   Perhaps they should have been called "abstraction neurons"

Competences need to be assembled in an architecture

How can we put everything together? We need to adopt the design stance (Dennett 1978) and make significant use of present and future concepts from information engineering and science. That will reveal a logical topography underlying the logical geography of concepts currently in use, pointing at the possibility of new deeper conceptualisations, as happened to ordinary concepts of kinds of stuff, following discoveries about the architecture of matter.

> See `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/logical-geography.html`

The Birmingham Cognition and Affect project has produced a draft high level classification of types of mechanisms, requiring many further subdivisions (the CogAff schema):

Requirements for subsystems can refer to

- Types of information used (ontology used: processes, events, objects, relations, causes, functions, affordances, meta-semantic....)
- Forms of representation (continuous, discrete, Fregean, diagrammatic, distributed, dynamical, compiled, interpreted...)
- Uses of information (controlling, modulating, describing, planning, executing, teaching, questioning, instructing, communicating...)
- Types of mechanism (many examples have already been explored – there are probably more to be discovered...).
- Ways of putting things together in an architecture or sub-architecture, dynamically, statically.

Perception	Central Processing	Action
	Meta-management (reflective processes) (newest)	
	Deliberative reasoning ("what if" mechanisms) (older)	
	Reactive mechanisms (oldest)	

In different organisms or machines, boxes contain different mechanisms, with different functions and connectivity, with or without various forms of learning. In some, the architecture grows itself after birth.

# A special case of the schema: H-CogAff

H-CogAff specifies a sub-class of human-like architectures within the generic "CogAff" schema. ("H" stands for "Human")

This is a sketchy indication of some of the required subsystems and how they are connected. Note the implication that both vision and action subsystems have several different (concurrently active) layers of functionality related to the different central layers/mechanisms.

Where could this come from?

Different trajectories for different layers:
- evolutionary,
  precocial competences from the genome
- developmental,
  altricial competences and architectures built while interacting with the environment
- adaptive changes, (small adjustments)
- skills compiled through repetition
- social learning, including changing personae...
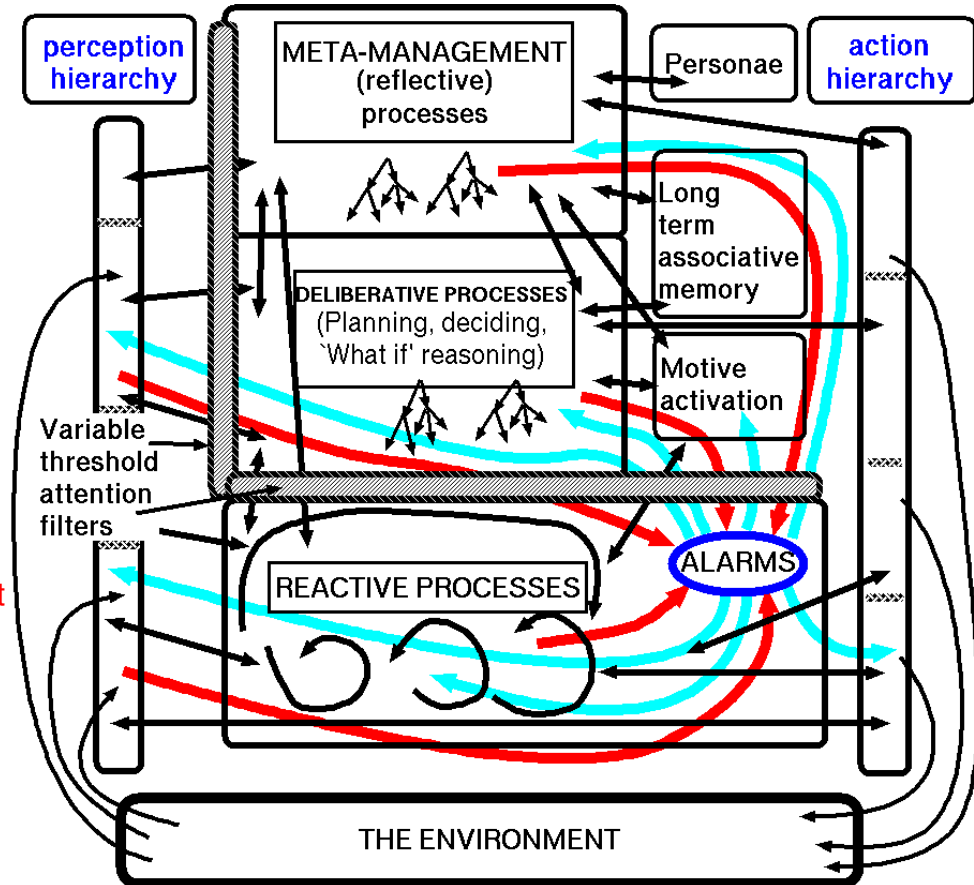
Much work remains to be done.

Kantian causal understanding and reasoning probably cannot occur in the reactive layers. Why not? Different variants may occur in deliberative and metamanagement layers.

For more details, see the presentations on architectures here: http://www.cs.bham.ac.uk/research/cogaff/talks/

Virtual machine events can also be causes

We now know that something like this can exist in computers, though there are people who dispute the causal links marked with '?'s.

Events in virtual machines, like inserting a character in a line, or a spelling checker correcting a typing mistake, or a chess program deciding to attack your queen are like mental events in many ways: e.g. such events have no physical location, cannot be observed using physical sensors, and are defined by concepts not definable in terms
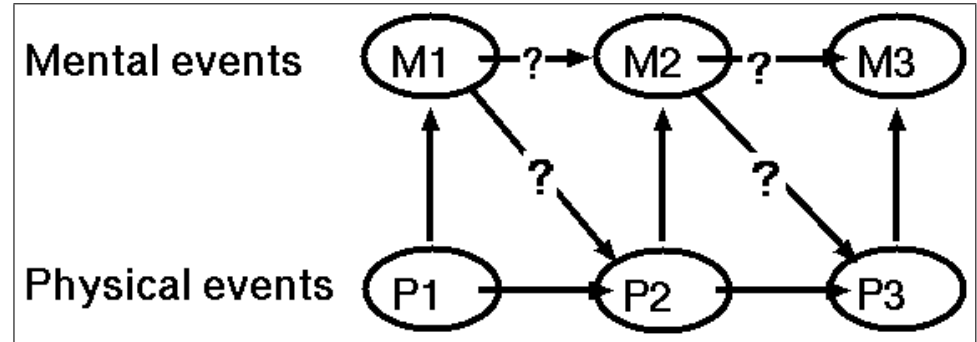
of the concepts of the physical sciences. Compare mental events.

Understanding all that is important for scientists (and philosophers) trying to understand minds, brains, and what they do and how they do it.

- People who challenge the links marked "?" (especially downward links) regard events in minds and computer virtual machines as epiphenomenal – i.e. as effects but not causes.
- This denial is based on a false theory of causation (e.g. if it is thought of as a kind of "conserved fluid" flowing through the universe).
- If instead we regard causal talk as merely a way of expressing which of a (rather complex) set of counterfactual conditionals are true and which false, then virtual machine events can be causes.
- (But that is not the main topic of this talk.)

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

Biological bootstrapping mechanisms

- There are some types of animals whose needs cannot be served by genetically determined (preconfigured) competences (using pre-designed architectures, forms of representation, ontologies, mechanisms, and stores of information about how to act so as to meet biological needs)

    why not?

- Evolution seems to have 'discovered' that it is possible instead to provide a powerful meta-level bootstrapping mechanism for 'meta-configured' species:

    - a mechanism without specific information about things that exist in the environment (apart from very general features such as that it includes spatio-temporal structures and processes, causal connections, and opportunities to act and learn, and that the neonate has a body that is immersed in that environment)

    - but with specific information about things to try doing, things to observe things to store

    - and with specific information about how to combine the things done and records of things perceived into ever larger and more complex reusable structures,

    - sometimes extending its own architecture in the process (e.g. in order to cope with a substantial extension to its ontology)

    - And including a continually extendable ability to run simulations that can be used for planning, predicting and reasoning.

So there are preconfigured and metaconfigured species, or, to be precise species with different mixtures of preconfigured and metaconfigured competences.

HYPOTHESIS

- In nature, fluid, flexible, metaconfigured, cognitive development (using particular sorts of architectures, mechanisms and forms of representation), is generally found only in species that biologists call 'altricial' – i.e. born/hatched under-developed and cognitively incompetent

    However, (a) the converse does not follow, and (b) the link is contingent:
    Are elephants exceptions?

- This underdevelopment and incompetence at birth may not be necessarily a feature of metaconfigured artificial systems with flexible cognitive development – perhaps some machines, or animals on some other planet, can be 'born' fully formed and fairly competent as well as possessing the competence to learn qualitatively new things by other means than slow statistics gathering.

- Nevertheless there may be design features that are required by both artificial and natural rapid and flexible learners, capable of spontaneously developing new ontologies and new combinations of old competences – e.g. if brain development has to be staggered or 'cascaded', then at birth infants are more likely to be incompetent.

    Compare cascade correlation neural network systems.

- We need to understand the design principles if we wish to develop machines capable of human-like understanding of the environment and rapid, flexible cognitive development.

- There can different competences in the same animal or robot – some more rigid (precocial, genetically determined) some more flexible (derived creatively from exploration and play).

- We need to understand relations between environmental and task constraints that favour different combinations of pre-configured and metaconfigured development.

Conjecture

We conjecture that rapid learning in altricial species depends on mechanisms making the metaconfigured learner spontaneously attempt many things without requiring a teacher.

Goals are generated in various ways by what is perceived.

> E.g. an infant that that cannot sit up or walk catching sight of the edge of the blanket may acquire the goal of grasping the edge, even if that means rolling over sideways. (See the movie!)

This depends crucially on discretisation (chunking) of continuous domains, to provide ontological and representational units that are capable of being combined in ever more complex discrete structures.

> Learn the easy things first, and some hard things become easy.
> It is nearly impossible to learn anything that is hard to learn.
> Oliver Selfridge: AI Magazine
> The Gardens of Learning: A Vision for AI
> 14(2): Summer 1993, 36-48

These ideas need to be expanded to include mechanisms that support substantive (i.e. non-definitional) ontology-extension. (I.e. abduction with introduction of new undefined symbols.)

How to control the search for such extensions is a major problem: partly solved by evolution's meta-configured competences.

If all this is correct it also undermines symbol-grounding theory and the sensory-motor contingencies theory of cognition.

How do you 'tell by looking'?

The examples of understanding (Kantian) deterministic causation in gears, links, shampoo containers, etc. presupposed that we sometimes can understand propagation of changes through changing structural relationships.

How it is done is far from clear, and it is far from clear how to implement such things in artificial systems.

The answer may be closely related to a theory of visual perception, according to which seeing involves running a collection of simulations at different levels of abstraction, partly, but not entirely, driven by the visual data.

Summary available here:
    http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505

- The simulation that you do makes use of not just perceived shape, but also unperceived constraints: rigidity and impenetrability.

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

KANT'S EXAMPLE: 7 + 5 = 12

Kant claimed that learning that 7 + 5 = 12 involved acquiring *synthetic* (i.e. not just definitionally true) information that was also not *empirical*. His idea may have been based on something like this simulation theory, as follows:

It is obvious that the equivalence below is preserved if you spatially rearrange the blobs within their groups:

```
ooo         o       oooo
ooo   +     o   =   oooo
o          ooo      oooo
```

Or is it?
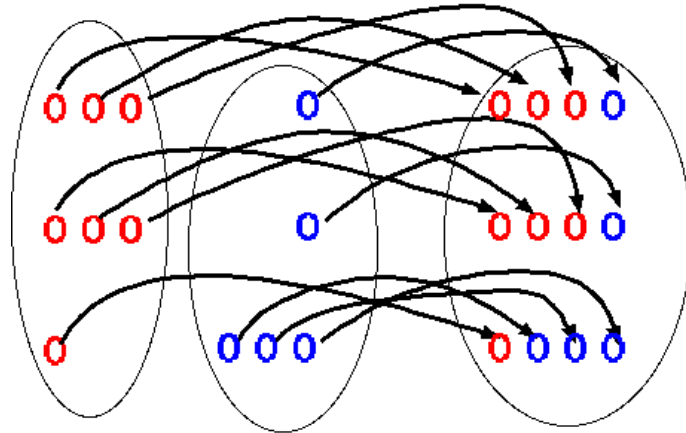How can it be obvious?
Can you *see* such a general fact?
How?

What sort of equivalence are we talking about?

I.e. what does "=" mean here?

Obviously we have to grasp the notion of a "one to one mapping".

That can be defined logically, but the idea can also be understood by people who do not yet grasp the logical apparatus required to define the notion of a bijection.
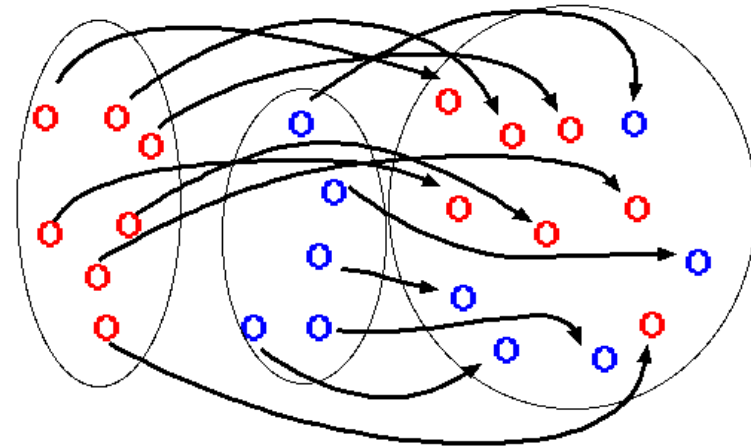
SEEING that 7 + 5 = 12

Join up corresponding items with imaginary strings.

Then rearrange the items, leaving the strings attached.

Is it 'obvious' that the correspondence defined by the strings will be preserved even if the strings get tangled by the rearrangement?

Is it 'obvious' that the same mode of reasoning will also work for other additions, e.g.

777 + 555 = 1332

Humans seem to have a 'meta-level' capability that enables us to understand why the answer is 'yes'. This depends on having a model of how our model works.

In the late 19th Century that was extended even to infinite sets

(e.g. by G. Cantor http://en.wikipedia.org/wiki/Georg_Cantor).

That's part of what we have to explain – on another occasion.

Different kinds of learning

- I am not claiming that learning about structure-based causation, or about mathematics, happens only in one way.

- Many competences that are first acquired and used on the basis of explicit reasoning mechanisms, are later used repeatedly:

    That causes another part of the relevant sub-system, a 'reactive' layer, to get trained to do the same task by going automatically from task or problem (plus context) to solution, using stored associations and very fast pattern-matching, instead of working out the required behaviour.

- This can allow tasks using highly trained subsystems to run in parallel, while the deliberative structure-manipulating creative learning subsystem does something else.

    There are many examples, some physical (e.g. learning to play musical scales fast or learning to ride a bicycle or drive a car), some perceptual, e.g. learning to read fast, and some mental, such as finding out numerical facts and then memorising them so that they are instantly available.

- Much learning of language seems to have the two strands: structure based explicit and relatively slow on the one hand and fast and fluent on the other.

    The latter fools some researchers into thinking it's all statistical.

- Many competences are divided into a subsystem that does things in a slow rule-based way and a subsystem that is trained to work fast and fluently, but is much more constrained in what it can do.

- Thus we should never ask 'How do humans do X?', for there may be different ways humans do X (walk, talk, sing, plan, see, think, learn, reason, predict, explain do calculations, ...) using different subsystems at different times.

    Future robots too!

Some Caveats

- I have emphasised the growth of understanding of the environment as based on a Kantian notion of causation – but only for some animal species, and some general-purpose robots, and some of their knowledge.

- Such understanding accounts for many of the most distinctive features of human life –
  and many causes of death also, when we act on incomplete or erroneous theories.

- Alongside growth of insight into spatial structures and how physical things work, a child also gradually bootstraps theories about how minds work, its own and others.
  Child science includes psychology (based on growing new meta-semantic competences).
  Likewise future baby-robots.

- However, I am not claiming that all or even most information about causation is based on explanatory knowledge about the underlying structures:
  Most of what a child learns about itself is Humean, including how to control its own movements, then later how various parts of its mind work, and how other minds work.

- Much self-knowledge, about body and mind, is incomplete, and liable to errors that persist in adult life (including religion and superstition), depending on how good the genetically determined and subsequently developed learning mechanisms are – and how far the understanding and teaching of science and engineering have progressed in the culture: often not as far as is assumed.

'Know thyself' Socrates is reputed to have said.
But understanding what is probably the most complex machine on earth, including many coexisting, interacting virtual machines within it, is easier said than done.

A Shared Research Agenda

The 'MeetingOfMinds' workshop presented knowledge that psychologists have that roboticists lack and vice versa, whereas I am concentrating on common gaps in knowledge requiring deep future collaboration.

- I am not simply recommending
  - producing lists of things discovered by psychologists, ethologists and neuroscientists that roboticists could learn from.
  - producing lists of mechanisms, algorithms, forms of representation, architectures, hardware-software designs, analyses of tradeoffs, etc. that AI researchers have investigated that psychologists, ethologists and neuroscientists could learn from.

- I have concentrated on deep common holes in knowledge, especially our understanding of varieties of perception, action, learning and causal understanding in a rich 3-D environment: where collaboration may accelerate progress.

- A possible outcome of the meeting would be creation of a shared research agenda, including further specification of the problems, summarising what is and is not known, and identifying a collection of worthwhile trajectories for future work.

- A part of this could be producing a partially ordered network of scenarios demonstrating various kinds of competence in humans and other animals and looking for empirically supported normal and abnormal trajectories (evolutionary and developmental) through the network, including unobvious intermediate cases, while trying to explain those trajectories by building working models of robots making different kinds of progress in different parts of the network.

- A major step would be producing a lightly annotated growing collection of videos demonstrating successes and failures in human infants, toddlers and other animals, and critically annotated videos showing the changing state of the art in robotics, AI vision, learning, etc.

Compare the euCognition Research Roadmap project: `http://www.eucognition.org/wiki/`

Some pointers – a tiny subset

Some papers and presentations written with Jackie Chappell.

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0502`
The Altricial-Precocial Spectrum for Robots (IJCAI 2005)

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609`
Natural and artificial meta-configured altricial information-processing systems (IJUC 2007)

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0702`
Computational Cognitive Epigenetics (BBS 2007). (Commentary on Jablonka and Lamb, cited previously.)

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#wonac07`
Slides for workshop on Natural and Artificial Cognition (NSF/EU Oxford, June 2007)

Interesting stuff on the internet about the nature of mathematics

`http://www.maths.manchester.ac.uk/~avb/micromathematics/`

`http://golem.ph.utexas.edu/category/2006/10/knowledge_of_the_reasoned_fact.html#c005571`
On why if a complex function of one variable has a first derivative all its derivatives must exist.

`http://www.kyb.mpg.de/~dcorfield`

Alexandre V. Borovik, *Mathematics under the Microscope: Notes on Cognitive Aspects of Mathematical Practice*
`http://www.maths.manchester.ac.uk/~avb/micromathematics/downloads`

Relevant work by Magnani discovered while I was expanding these slides

`http://www.unipv.it/webphilos_lab`

Magnani, L.:
Multimodal abduction. External semiotic anchors and hybrid representations.
Logic Journal of the IGPL 14(1) (2006) 107-136
`http://www.unipv.it/webphilos_lab/papers/magnani3.pdf`

Magnani, L.:
Animal Abduction - From Mindless Organisms to Artifactual Mediators,
Studies in Computational Intelligence (SCI) 64, 3-37 (2007)
Springer Verlag
Berlin Heidelberg 2007
`http://www.unipv.it/webphilos_lab/papers/magnani1.pdf`