Acquiring Orthogonal Recombinable Competences

Poster for COGSYS II Conference, Nijmegen, April 2006

Aaron Sloman, Birmingham CoSy Project Team, and Jackie Chappell,
University of Birmingham, United Kingdom

A child or baby robot that has to manipulate 3-D objects in its environment would face a combinatorial explosion if all possible situations have to be learnt about separately. This could take evolutionary time-scales.

It is conjectured that humans and some other altricial species instead use innate mechanisms for decomposing situations into components that can be explicitly learnt about, and stored in such a way that the competence can be re-used in combination with other learnt competences, in perceiving novel situations and performing novel actions.

That includes learning about kinds of surface fragments (e.g. varieties of curvature and surface discontinuities), kinds of surface properties (e.g. texture, hardness, etc.), kinds of material (rigid, flexible in different ways), kinds of objects composed of materials and shapes, kinds of relationships, kinds of changes in relationships, kinds of causal connections between changes.

These need to be represented in a manner that is independent of precise sense-data when they are perceived, or sensorimotor contingencies, so that knowledge about them can be used in planning future actions, thinking about the past, and comparing actions using different hands, or hands or mouth in different positions. This implies a use of 'objective' representations (e.g. of 3-D structure) which can then also be used in perceiving 'vicarious' affordances (for others).

An implication is that mirror neurons should have been called 'abstraction neurons'. There are many other implications, for robotics, psychology and neuroscience.

The Child as Scientist:

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

Video available at

http://www.jonathans.me.uk/josh/movies/josh23_0040.mpg

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 ???

The baby seems to experiment with his hands, legs, the spoon, the yogurt and the carpet. Perhaps he is learning what causes what? What does that mean?

Kinds of Causation: 1 (Humean)

Two gear wheels attached to a box with hidden contents.

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

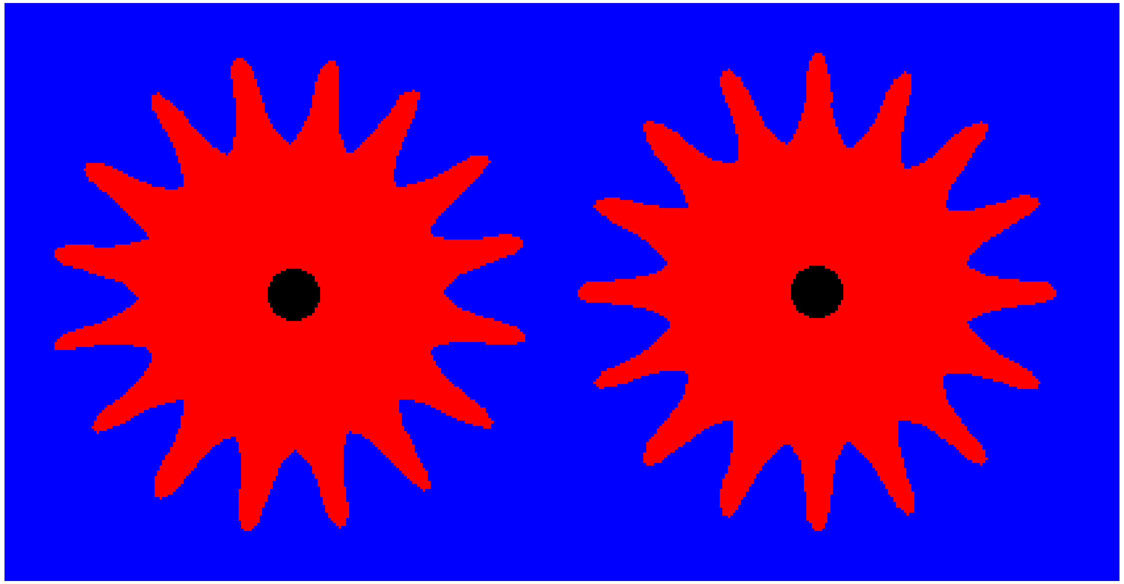

- **You can tell by experimenting: you may or may not discover a correlation – depending on what is inside the box.**
- **In more complex cases there might be a knob or lever on the box, and you might discover that which way the second wheel rotates depends on the position of the knob or lever. (Compare 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.**

Kinds of Causation: 2 (Kantian)

Two more gear wheels:

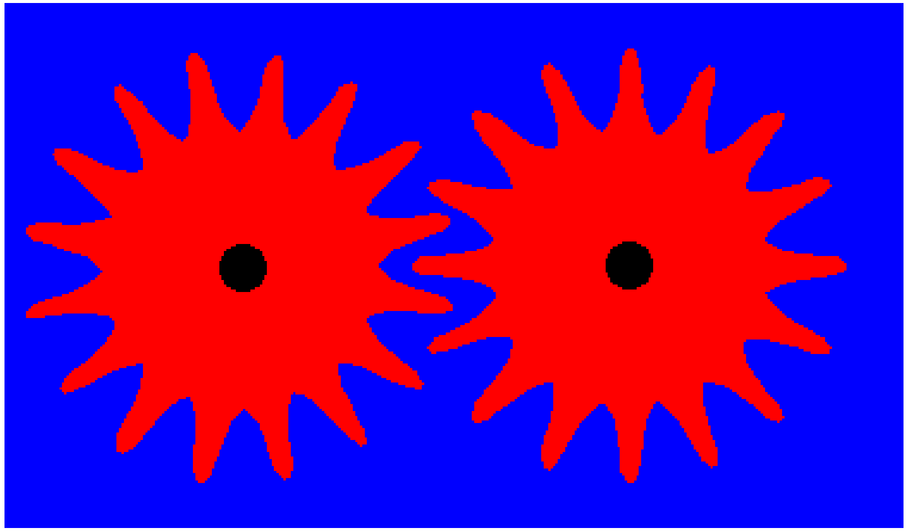

You (and some children) can tell, by looking, how rotation of one wheel will affect the other.

How? You can simulate rotations and observe the consequences.

What you can see includes this:
As a tooth near the centre of the picture moves up or down it will come into contact with a tooth from the other wheel. If 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.

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

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

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

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

The ability to do causal reasoning in different domains has to be learnt.

The ability to work out consequences in novel contexts requires learning to build simulations with appropriate structures, appropriate permitted changes, and appropriate constraints.

What is appropriate depends on what is being simulated: simulating the rotation of a rigid gear wheel (e.g. one made of steel) is not the same as simulating the rotation of something soft and malleable, e.g. putty or plasticine.

Appropriate constraints ensure the right counterfactual conditionals are true as the simulation runs.

The detailed representational, algorithmic, mechanistic and architectural requirements to support such learning, and the growth of the ontology involved, require much deeper analysis than I can give at present.

Part of the point of the CoSy 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.

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.

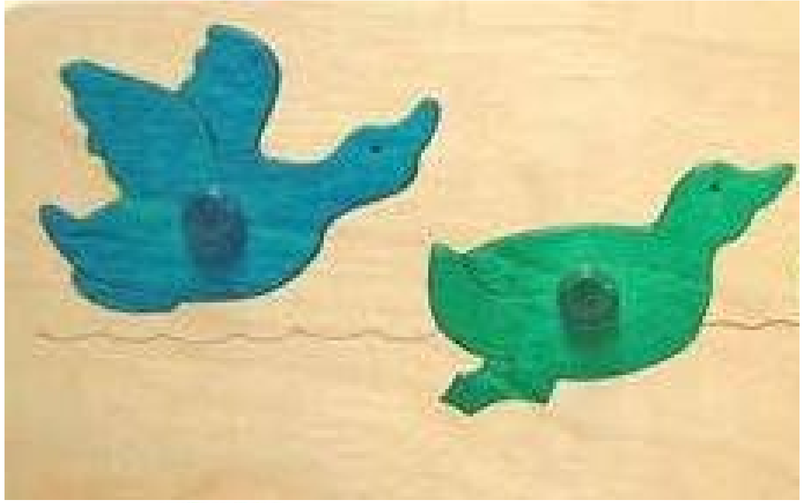

A child learns that it can lift a piece out of its recess, and generates a goal to put it back, either because it sees the task being done by others or because of an implicit assumption of reversibility. At first, even when the child has learnt which piece belongs in which recess there is no understanding 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.

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 the child's mental simulation mechanism. HOW?

A succession of stages

- **The process of extending competence is not continuous (like growing taller or stronger):**
- **The child has to learn about**
  - **distinct new kinds of objects, properties, relations, process structures, e.g. for rigid objects, flexible objects, stretchable objects, liquids, sand, treacle, plasticine, pieces of string, sheets of paper, construction kit components in Lego, Meccano, Tinkertoy, electronic kits...**
  - **new forms of representation, new kinds of transformations, new constraints on transformations, new applications of recent acquisitions.**
- **The word 'stage' can mislead: 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.**
- **I don't know how many different things of this sort have to be learnt, but it is easy to come up with hundreds of significantly different examples.**
- **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.**

  CONJECTURE:

  in the first five years a child learns to run at least 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.

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 I had ironed the handkerchief 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 either case a square can be folded to make a rectangle.

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.

On the other hand 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.

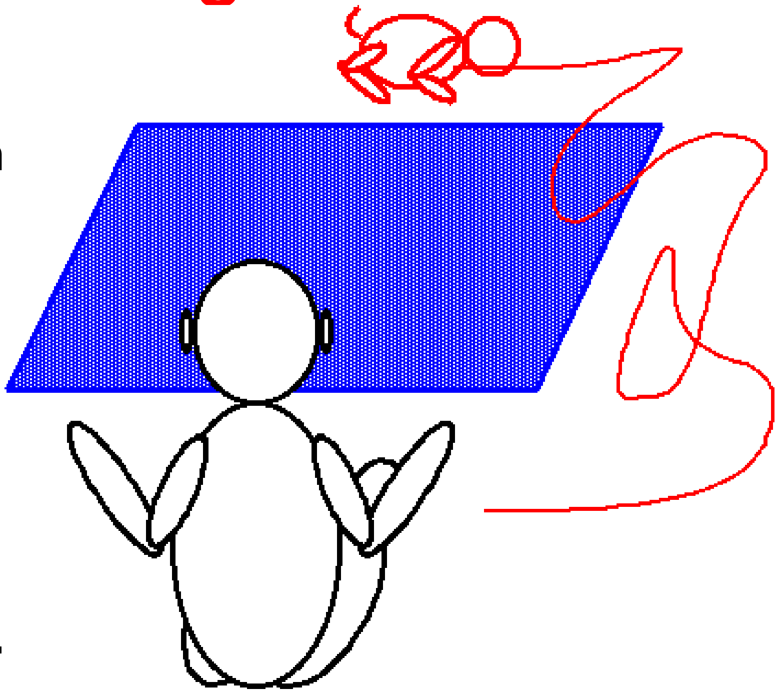

However the ontology of strings is a bag of worms, even before knots turn up.

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.

Sensory-motor vs action-consequence contingencies

Two evolutionary ‘gestalt switches’?

The preceding discussion implies that during biological evolution there was a switch (perhaps more than once) from

> **insect-like understanding of the environment in terms of sensory-motor contingencies linking internal motor signals and internal sensor states (subject to prior conditions),**

to

> **a more ‘objective’ (allocentric?) understanding of the environment in terms of action-consequence contingencies linking changes in the environment to consequences in the environment,**

followed by

> **a further development that allowed a generative representation of the principles underlying those contingencies, so that novel examples could be predicted and understood, instead of everything having to be based on statistical extrapolation.**

A major driver for this development could be evolution of body parts other than the mouth that could manipulate objects and be seen to do so.

> **However the cognitive developments were not inevitable consequences: e.g. crabs that use their claws to manipulate food do not necessarily have the generative competence.**

(I was told by a biologist at the conference that the above was unfair to insects)

Biological bootstrapping mechanisms

- **There are some species 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)**
  - **with specific information about types of: 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.

Biological Nativism: Altricial/Precocial tradeoffs

- Evolution 'discovered' that for certain species which need to adapt relatively quickly to changing environmental pressures and which perform cognitively demanding tasks as adults, a kind of Kantian learning mechanism is possible that allows much faster and richer learning than is possible in systems that merely adjust probabilities on the basis of observed evidence (statistical data).
- The latter species, with more or less sophisticated forms of the Kantian mechanism, learn a great deal about the environment after birth and and in some cases are able rapidly to develop capabilities none of their ancestors had (like young children playing with computer games).
- We conjecture that this uses an information-processing architecture which starts off with a collection of primitive perceptual and action competences, and also with a mechanism for extending those competences by 'syntactic' composition, as a result of play and exploration, which is done for its own sake, not to meet other biological needs (food, protection from hurt, warmth, etc.)
- The meta-level features of the mechanism and the initial competences are genetically determined, but the kinds of composite competences that are built are largely a function of the environment.
- This requires forms of learning that are not simply adjustments of probabilities, but involve continual creation of new useful structures.

Summary so far:

There is an important sub-class of animals in which competences are not all pre-configured, whose development makes use of:

- **Genetically determined primitive actions, perceptual capabilities and representations,**
- **Genetically determined play/exploration mechanisms driving learning that extends those actions, etc., using abilities to chunk, recombine and store**
  - **new more complex action fragments**
  - **new more complex perceptual structures**
  - **new more complex goals**
- **Creating new ontologies, theories, competences (cognitive and behavioural),**
  - **i.e. new more complex thinking resources,**
- **Thereby extending abilities to search in a space built on larger chunks: solving ever more complex problems quickly.**
  - **(unlike most statistical forms of learning)**
- **Humans are able to apply this mechanism to itself – producing new forms of self-awareness and new forms of self-understanding, including mathematical knowledge.**

For AI systems this will require us to discover new architectures and learning mechanisms.

A spectrum of competences

- Every organism is a mixture of both kinds of capabilities:

  pre-configured — constructed

- Not all of the first kind are manifested at birth/hatching – many are ‘time-bombs’.
- Architectures for altricial species can do many things that are not directly biologically useful:

  including (possibly dangerous) exploration of a space of possibilities.

- Architectures can change over time.
- Altricial architectures are virtual machines that grow themselves.

  But we have over-simple ideas about how: e.g. the notion of a knowledge-free, general-purpose learning system is current favourite, but inadequate mechanism.

See our (Sloman & Chappell) IJCAI paper
http://www.cs.bham.ac.uk/research/cogaff/05.html#200502

and the H-CogAff architecture described on the Cognition and Affect web site:
http://www.cs.bham.ac.uk/research/cogaff/

Implications for theories of meaning

The existence of precocial species refutes 'symbol-grounding' theory

(Otherwise known as 'concept empiricism' – the theory that all meaning has to be derived by processes of abstraction from sensory experiences, which is clearly not required for precocial species that are competent at birth).

In our IJCAI paper we distinguish two sources of meaning

- **the structure of a theory in which 'undefined terms' occur**
  (where the structure limits the class of possible models/interpretations)
- **links to sensing and acting**
  (where the links – e.g. Carnapian 'meaning postulates' further reduce the set of possible interpretations, tethering the interpretation – though there is always residual indeterminacy.)

The second picture seems to represent how scientific theories get their meaning, so why not toddler theories?

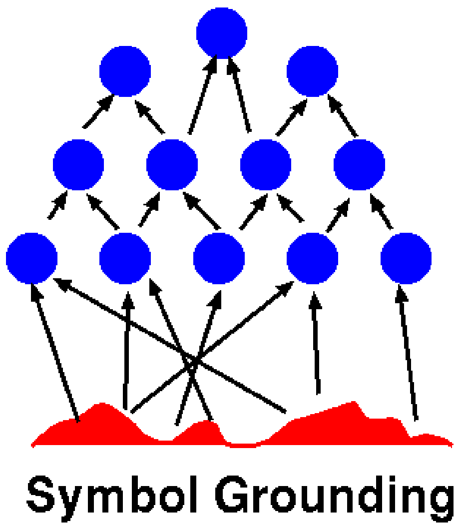

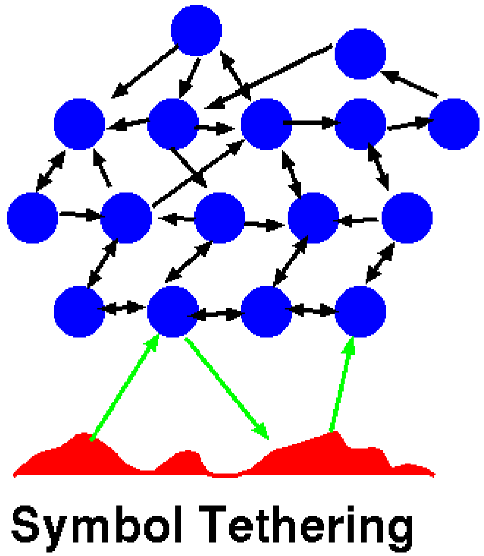

additional points

- We need to find out how many different kinds of simulative capabilities a child, a chimp, a nest-building bird, a domestic robot, needs to acquire.
- We need to understand what sorts of forms of representation, mechanisms and architectures, can produce those developments.
- The process can involve creation of new ontologies and new forms of representation.
- There will not be a simple step: understanding causation
- Many different kinds of cognitive competence relevant to understanding different kinds of structures and processes grow during our life time.
- Different people grow different subsets (why?)
- Scientific research is just an extension of this – though too many scientists restrict their research to accumulation of correlations (like learning in precocial species?).
- When the ability we are discussing is applied to itself we get activities like mathematics and philosophy.

Conclusion

- I have been emphasizing the growth of understanding of the environment as based on a Kantian notion of causation – but only for some altricial species.
- This 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.
- However I am not claiming that all or even most of our information about causation is based on explanatory knowledge about the underlying structures.
- In particular, most of what a child learns about itself is Humean, including how to control its movements, then later much of how its mind works.
- Much self-knowledge, about body and mind, is incomplete, and liable to error.
- Alongside growth of insight into how physical things work a child also gradually bootstraps theories about how minds work, its own and others – child science includes psychology as well as mechanics and physics.

  Both can produce errors (including religion and superstition) that persist in adult life. The errors will depend 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.

'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.

See also: http://www.cs.bham.ac.uk/research/cogaff/talks/

Further Reading

This poster presentation was produced by starting from a much longer presentation at a conference and leaving out most of the slides.

The longer presentation can be found at

http://www.cs.bham.ac.uk/research/cogaff/talks/aisb06-ortho.pdf

A web site exploring some of the issues in more detail is here

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601
Orthogonal Recombinable Competences Acquired by Altricial Species (HTML)

A (possibly) new theory of vision related to all this is presented in

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505 (PDF)

A discussion of ways of conceiving of and learning about causality

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506 (PDF)
Two views of child as scientist: Humean and Kantian

A paper co-authored by Jackie Chappell, presented at IJCAI'05

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