

Presentation at Workshop on Computational Approaches to
Representation Change During Learning and Development
AAAI'07 Fall symposium, Washington 10th Nov 2007

May be revised

Diversity of Developmental Trajectories in Natural and Artificial Intelligence

Aaron Sloman

University of Birmingham, UK

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

In collaboration with Jackie Chappell

These slides will be in my 'talks' directory:

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

The position paper for the meeting is available here

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

Apologies

I apologise

- For slides that are too cluttered: I write my slides so that they can be read by people who did not attend the presentation.

So please ignore what's on the screen unless I draw attention to something.

My presentation uses linux and latex

Different Goals for AI Research

My primary aim is not to produce intelligent machines

(except as experiments testing the theories and models).

This is a presentation from the viewpoint of a researcher trying primarily to understand and model aspects of **natural** intelligence, in its many and diverse forms,

... but not only human intelligence: also microbes, insects, birds, deer, hunting mammals, primates, landmark-using ants, portia spiders, Betty the hook-making crow,

However, if we ever do understand the products of biological evolution well enough to produce convincing working models, that may help us to become more effective at producing intelligent artifacts to meet many kinds of practical needs.

But producing something that works is not enough: we need to understand:

- What features of the design make it suited to the requirements (environment, niche);
- How different requirements might have required different designs;
- How different designs might have met the requirements in different ways;
- What the tradeoffs are.

Understanding goes beyond modelling.

Understanding Goes Beyond Modelling

Barbara Webb on relevance of models to Biology:

“.... unless, at some point, an animat model can be shown to account for or match some real and significant biological data, it is hard to argue convincingly that it is relevant to biology. But this is not the same as requiring the implemented model to be a detailed, structurally accurate representation of the mechanism producing this behaviour in the biological system.”

Webb, B. (2007) Animals versus animats: or why not the real iguana? (to appear in) *Theoria et Historia Scientiarum*

http://homepages.inf.ed.ac.uk/bwebb/publications/Webb_THS-ALife.pdf

I'll try to explain a different way of doing biologically relevant AI research.

I am trying to do something different:

Instead of trying to model any particular organism

(a perfectly worthy objective, though not mine – except as a means to a broader objective)

I am trying:

- to understand how the tasks of modelling different sorts of organisms (and also possible machines) are related;
- hoping to produce a unifying theoretical framework, relating different sorts of models, of different sorts of animals and machines,
- by analysing how demands and opportunities in the environment vary,
- focusing on whole organisms/machines rather than on component mechanisms (vision, motor control, language, learning, etc.)
- inspired by demands that changing environments place on evolutionary and developmental trajectories.
- looking at “inner logics” of some of the trajectories,
e.g. trying to understand why design X should precede design Y, and why new pressures (changes in climate, predators, food supply, etc.) are not coped with as well by design X as by a new design Y.
NB: the constraints probably produce partial not total orderings: different routes to similar ends.

This treats evolutionary and developmental trajectories not just as brute empirical facts, but as partly explainable by the logical relationships between sets of requirements.

How to carve up the research space

One possibility

- Everyone chooses some particular practical goal (or modelling goal) and tries to design an intelligent machine to achieve that goal, or works on a particular type of mechanism trying to improve its performance on an agreed set of tests.

Another possibility

- Some people try to work on understanding the space of possible goals, and requirements, and the space of possible designs for meeting those requirements.

How can we carve up such spaces?

- study different sorts of **competences** and **behaviours**, and types of **mechanisms** capable of supporting them; and use competence-based metrics for comparing designs (This is how a great deal of AI research now works.)
- study different sorts of **complete** systems and how to model or implement them
- identify different kinds of **environment** and study requirements for coping within them
- identify and compare different sorts of evolutionary and developmental **trajectories** and
 - the mechanisms involved in supporting the trajectories
 - the varying relationships between designs linked by trajectories
 - perhaps focusing on particular **neighbourhoods** in design space or niche space

Dynamics of Linked trajectories

Movements along trajectories in design space cause, and are caused by motion along trajectories in niche space:

Change in design of species A changes requirements for species B, which causes changes in design of species B, which changes requirements for A, etc.: there are networks of such interactions.

This obviously applies to e-trajectories, and less obviously to i-trajectories

Niches for an unborn foetus, for a newborn infant, a schoolchild, a parent, a professor, etc. are different – though the same human can be in them all at different times.

Moreover, an individual can instantiate more than one design, satisfying more than one niche: e.g. switching between being

- protector and provider,
or
- parent and professor

To cope with development of multi-functional designs we can include *composite niches* in niche space, just as there are composite designs in design space.

Composite niches lead to composite designs and vice versa.

New mathematics needed?

The history of the biosphere involves multiple interacting e-trajectories for designs and niches, with many interacting feedback loops.

Previous slides suggest a need for new mathematics linking all the many complex feedback loops involved at different levels of abstraction between niches and interacting design instances.

Biological evolution:

As more and more complex organisms evolved, their i-trajectories became longer and more diverse.

Much later came s-trajectories and c-trajectories,

Maybe there will soon be r-trajectories (genetic engineering?)

Many questions: e.g. why are there so few “intelligent” species or individuals?

(Count species, individuals or biomass.)

Answer: the most sophisticated brains have to be near the tip of food pyramid.

Under what conditions does the (expensive) transition to deliberative capabilities pay off, compared with other design options?

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604>

Requirements for a Fully Deliberative Architecture

Are those conditions very rare?

How many intermediate cases are there between purely reactive and fully deliberative mechanisms?

How many different forms of representation did evolution produce, and why?

What led to development of a-modal non-sensorimotor ontologies?

Why do we get some apparently intelligent species (mammals, birds) that are **precocial** while others are **altricial**?

Are there some constraints that require more sophisticated adults to start from less sophisticated infants?

How i-trajectories start: tabula rasa or ...?

The fact that some species, the precocial species, start so competent provokes the question:

Why do other species, such as primates, hunting mammals and nest-building birds, start so helpless and incompetent?

Such species are labelled 'altricial'.

Do they start with no knowledge and only a very general learning mechanism, or do they have sophisticated but 'generic' knowledge or 'meta-knowledge'?

Evolution did not solve one problem with one solution:

(See papers and presentations by Chappell and Sloman)

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.

Compare John McCarthy: The Well-Designed Child

<http://www-formal.stanford.edu/jmc/child.html>

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

.....

Instead of building babies as Cartesian philosophers taking nothing but their sensations for granted, evolution produced babies with innate prejudices that correspond to facts about the world and babies' positions in it. Learning starts from these prejudices. What is the world like, and what are these instinctive prejudices?

“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.”

<http://www-formal.stanford.edu/jmc/child.html>

Biological facts support McCarthy

Most animals start life with most of the competences they need apart from some fine tuning – e.g. deer that run with the herd soon after birth, chickens(?), invertebrates, etc.

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..

Blooming buzzing confusion

If John McCarthy is right, William James got it badly wrong.

Of course he did identify a problem.

But it is not the problem solved by baby mammals or birds.

It may be a problem solved by some other organisms, possibly long dead organisms.

Compare the book by Eleanor Gibson and Anne Pick
An Ecological Approach to Perceptual Learning and Development,
2000

Varieties Of Worlds To Learn In

One way to explore the issues is to consider different sorts of worlds in which organisms might

- behave
- develop
- evolve

We can see how different environments produce both

- Different constraints
- Different opportunities

(Gibson's notion of "affordances" for an individual can be generalised to **affordances for a species, or gene pool.**)

It's not only the environment:

As a species evolves, acquiring more complex body parts and more complex information processing capabilities, its constraints and opportunities change.

Changes in designs produce new requirements: new niches for subsystems

E.g. having an **articulated body** provides major new opportunities, and far more demanding requirements beyond mere mobility.

From chemical soups to skyscrapers

The AAAI'07 Symposium paper presented a few snapshots from a large space of possible environments, illustrating some of the ways in which the affordances in those environments vary.

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

The examples include:

- amorphous chemical soups
- soups with some spatial structure
- soups with additional impervious structures, e.g. walls, forming corridors and doorways
- objects that need to be assembled from parts, or disassembled to provide parts, etc.
- Structured animals with independently movable parts
- Animals that need to enhance themselves with tools as the things they build get larger, more complex or more sophisticated.
- Animals with complex virtual machines attempting to understand themselves
- Animals with social and competitive relationships with other animals.

This is not an exhaustive survey, just an illustration of a type of methodology not widely recognized as appropriate to AI

Soup dwellers:

A chemical soup, even with different chemicals needed at different times, does not give much scope for cognition

at most information about local gradients can be used.

If there is some other field, e.g. gravitation or illumination field-strengths may correlate with nutrients and 'noxients': that may allow controlled traversal based on gradients to optimise feeding.

If things are changing there may or may not be learnable patterns in the changes.

Varying soup-challenges:

If there are permanent obstacles separating regions with different chemicals, use of gradients may still suffice but more flexibility is possible if the geography can be discovered and used

E.g. maps of walls, corridors, rooms, with regions where different chemicals have greatest concentrations, etc.

NB: new architectures, forms of representation, mechanisms are needed to acquire, store and use such information.

Other possibilities, distribution of nutrients and noxients, competitors for the same nutrients, predators, mobile food (prey organisms)

NB: Learning can be done either by evolution or by individuals – with different costs and benefits

Noxious and walled soups:

Fixed, impenetrable barriers to motion can also add opportunities to benefit from the ability to learn and remember their locations.

The barriers impose rigid constraints on motion whereas the noxients impose soft constraints.

Not yet the problem Ben Kuipers discussed, because the sensors still record only local information whereas he was using laser range finders.

Atomic sensors:

If new types of substance (nutrients and noxients) are introduced or removed from time to time, then the problem for evolution is much harder.

It needs to allow the information-processing architecture and the contents of representations to be take different forms in different generations.

There are two main ways of doing this.

If the changes are slow enough then, as the soup-world changes, the fixed architecture of the organisms can slowly vary from generation to generation, and over time some sequences of variation may produce surviving sequences of organisms in the new environments.

If environmental changes occur faster, then **individual** organisms need to be able to learn about the new substances.

That requires sensors to have more generality:

e.g. by being able to detect 'lower-level' features that are common to different chemicals and finding out which combinations are good and bad, etc.

NB. It makes a big difference whether different types can be characterised simply by different vectors or whether **structural descriptions** are needed (e.g. for same atoms arranged differently).

From Moving to Manipulating:

Some environments make it possible for smaller fragments to be combined to form useful larger objects, and larger objects to be disassembled and fragments re-used.

This provides yet more challenges for evolution and learning.

This is also a requirement for digestion, metabolism, growth, repair of damage, immune responses – all of which use chemical mechanisms produced mostly by evolution, which operate in purely reactive ways in various parts of an organism's body)

Manipulation of external complex objects introduces new demands.

The effectors need to be more complex, so that objects or their parts can be moved together or apart or rearranged.

This requires the ability either to move one object to or away from another, including rotational movements when constructing 3-D objects (unless made from something like mud).

Some complex **concurrent** movements may be required.

These opportunities also substantially extend the requirements for information-processing.

Whereas previously the only kind of future to be considered in formulating goals or predicting consequences of actions was a future in which the organism's location (and possibly orientation) changed, manipulation involves far more varied changes than transformation of a fixed-size vector.

If planned objects can be more or less complex, that will require the ability to construct more or less complex representations of objects, instead of using only fixed-size, fixed-complexity vectors of measurements.

Exosomatic ontologies:

Exactly how these new structures should be represented is debatable.

An organism that can represent only relations between signal-patterns within its body (sensorimotor patterns and contingencies) uses only a 'somatic' ontology.

The ontology can be uni-modal or multi-modal.

If an animal or machine uses an ontology that refers to things that can exist or occur independently of how they are sensed or acted on by that agent, it uses an 'exosomatic' ontology.

Components of an exosomatic ontology require a type of semantic content that is not definable in terms of sensori-motor (somatic) concepts.

Such concepts cannot be accommodated within symbol-grounding theory.

However they can be represented by undefined symbols in a theory about the environment.

Their semantics may be partly implicitly defined by their role in the theory (model-based semantics), and partly by bridging rules (theory-tethering).

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

There are many advantages to exosomatic ontologies including their economy in representing what is common to many different processes (or actions) and their use in making predictions, forming plans, forming explanations, including hypothesised histories.

From Inheriting to Discovering:

If the kinds of object and kinds of manipulation required in a particular environment do not change much then it is possible for evolution to produce combinations of sensors, effectors, and information-processing mechanisms, including ontologies and forms of representation required.

Compare precocial species.

Evolution needs to be replaced by learning if the environment changes faster than evolved designs can, for instance

- by providing new sorts of materials,
- producing new configurations of objects,
- producing new climatic conditions, e.g. temperature,
- new environments found by migration
- new behaviours in rival species
 - (prey, or predators, or competitors for the same food or habitats)
- new arrivals from other locations

If these changes happen too fast, it will not be possible for evolution to hard-code all the required competences, even if the sensors and effectors provided are potentially very general.

Then some of the burden of accommodation needs to be shifted to learning.

Manipulation Changes What Can Be Learnt

In a 3-D world, as more and more complex objects are constructed by assembling available components, those complex objects themselves can be parts of still more complex objects.

So the fact that certain objects have been made makes it possible to discover that there are new kinds of objects that can be made in fairly small steps that could not previously be made in small steps.

As objects become larger the problems of manipulating them change, also requiring more complex combinations of actions to be created.

Likewise as actions are produced that assemble such objects, 'chunked' combinations of actions become available to be used as components in more complex actions

Sometimes the complexity involves doing longer sequences of things, and at other times it involves doing more things in parallel, using cooperative agents.

An important point on which much more research is needed is that spatial relationships between processes can cause complex interactions which may provide rich and complex positive and negative affordances.

This is relevant to understanding causation in 3-D structures and processes

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

The value of sequential chunking in reducing complex search spaces has been well known in AI for decades. I don't think chunking of concurrent spatially integrated actions has been studied much.

Ontology extension:

The ontology of the learner may also have to be extended to include new kinds of design, new kinds of tool, new kinds of construction process, new ways of collaborating with others, providing new contents for thinking, seeing, planning, and learning processes.

[Show video:

[Yogurt can be food for the mind as well as for the body in an 11 month old scientist. \]](#)

Watching young children playing with toys of various sorts shows that things that seem obvious to older children may be completely incomprehensible to younger ones for a while, such as why putting a puzzle piece in its recess and pressing does not make it return to its previous location.

[Show video of betty?]

Videos of the crow, Betty, spontaneously making hooks in several different ways, are available online at the Oxford zoology web site. Use a web search for “betty”, “crow” and “hook”.

Learning begets new needs:

By enabling the learner to produce novel structures, those learnt abilities themselves rapidly produce new opportunities (and sometimes threats – since new constructions can be dangerous as well as useful), possibly requiring another ‘layer’ of learning.

The learning capabilities that produced the early competences will not always suffice for producing the newer more complex ones.

It seems that somehow this was discovered by evolution.

What can be learnt by altricial individuals at various stages changes significantly.

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.

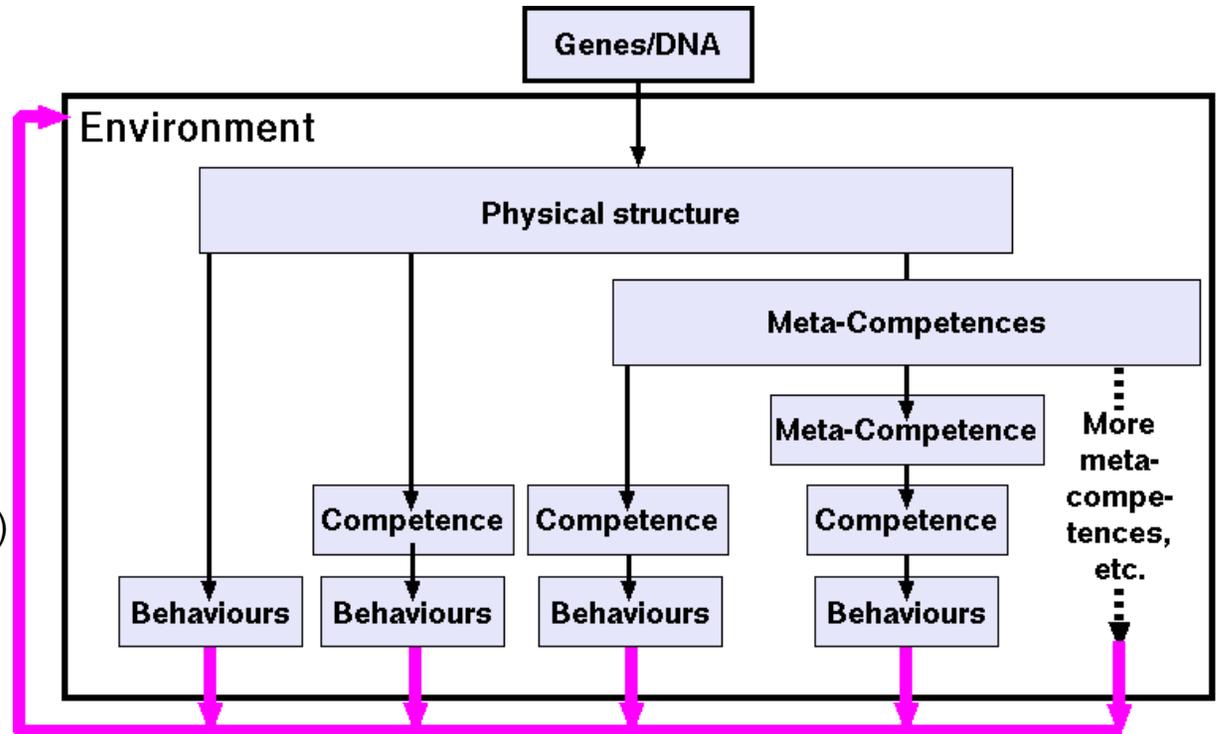
They occur towards the left.

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.



(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?

Innate Context-Specific Meta-knowledge:

If the previous arguments are correct, then some meta-competences that enable or facilitate the acquisition of new competences (or new meta-competences), far from being general learning algorithms, are specifically tailored to finding things out about restricted environments.

See later slides

Definitional vs Substantive Ontology-Extension

Patterns in sensory and motor signals define an immediately accessible ontology which is the only kind simplest organisms can use. (Compare Braitenberg's "vehicles".)

Exosomatic ontologies, e.g. referring to things in the environment that exist independently of being sensed or acted on require **substantive** ontology extension – a hard task, whether for e-trajectories or i-trajectories.

- This would be impossible if 'symbol-grounding' theory (concept empiricism) were true!
- **Definitional** (abbreviative) ontology extension uses new combinations of existing concepts (including relational concepts and logical connectives), to form new concepts that are found to be useful: they are stored for re-use and associated with a new label.

A more sophisticated type of abbreviative ontology uses **invariants found in patterns of change**.

E.g. if electric currents and voltages can be measured then a conductor may be found to exhibit an invariant ratio between current and voltage as current and voltage change: that is given a label "resistance".

Many new concepts are introduced that way in science, and AI researchers and cognitive scientists have attempted to generalise this to more sophisticated sensorimotor invariants (e.g. O'Regan).

- Abbreviative labels are heuristically useful but do not **extend** the expressive competence of the learner.
Like "macros" they sometimes help the search for new useful combinations of old concepts.
- Abbreviations pick out useful subsets from the very large space of possible concepts and possible laws already expressible in the learner's ontology.
- **Substantive ontology extensions** go further.

Substantive concept learning:

- In contrast, **substantive** concept learning produces new concepts that are not definable in terms of the initial set, allowing construction of new theories or hypotheses that were not previously expressible.
- Some previously known facts may later turn out to support or contradict those theories, and old puzzles may be explained by the new theories.
- Such cases are familiar from the history of science.
- Substantive ontology extension includes learning about different kinds of matter and properties of various materials that are not detectable using available sensors, for instance, solubility or atomic number of elements.
- What is learnt through the application of meta-competences includes what sort of ontology is useful in the environment, as well as which laws using that ontology work well for making predictions in the environment.
- But it is not a simple algorithmic process, for it involves controlling search in the vast space of ontology extensions.
- The search is controlled by introducing new concepts only when a current explanatory **theory** needs to be extended or modified because some plans, predictions and explanations turn out to be inconsistent with available evidence.
- Theory formation is called “abduction”, but AI abduction systems do not usually allow ontology extension (use of new undefined symbols in “axioms” added to a theory).

“Inward” ontology extension

An organism may need sophisticated forms of control that requires it to have a model of itself, as has often been proposed. (E.g. Owen Hollands, and our meta-management architectural layer).

This inward looking theory about the world will often be most usefully a theory of high level virtual machinery rather than the detailed physical mechanisms and processes.

For this the system will have to go beyond what internal physical sensors can sense: another example of substantive ontology extension.

Forms of representation for 'inner languages'

Previous work with biologist Jackie Chappell suggests that in some species the kinds of perceptual, planning, problem-solving, and plan execution competences that develop require sophisticated internal forms of representation that we call 'generalised-languages' abbreviated as 'g-languages', or GLs.

A GL has three features often assumed to exist only in human languages,

- structural variability (including Fregean and analogical formalisms)
- (context sensitive) compositional semantics
- goal directed manipulation

GLs

(a) Structural variability in a GL allows the construction of complex information structures with varying numbers of parts in varying relations.

This includes the use of lists, trees, and networks containing different sorts of items of information. It also includes distributed as well as localised and geometrical as well as logical forms of composition.

(b) Compositional semantics in a GL allows any information structure to occur as parts of several different larger information structures,

the information (meaning) expressed in the larger structure will be determined by (i) the structures of which it is composed, (ii) how they are organised, along with (iii) relevant contextual information.

Different notions of part, whole and composition to form complex information structures are possible.

An example of use of a GL would be Use of a GL is essential for 'fully deliberative' architectures (described in <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604>).

The parts in an intended configuration of objects could include parts of the animal, e.g. hands, feet, or jaws, which might be required in intermediate phases of the production of some new complex object. Animals building nests, peeling bananas, or using two rocks to crack goal state (e.g. termites?).

Alternatively such actions may depend on using a GL to represent a new goal which is then used to control actions constructing, disassembling or rearranging objects, using do next.

Very few types of animal can do this.

Deciding which ones can is hard. (Chappell/wonac)

Representation is always for an information user and constrained by that user's information-processing capabilities.

GLs precede communicate languages

GLs must have evolved for internal use prior to the evolution of human language – though their use for external communication probably accelerated their development.

Likewise GLs must be available for internal use prior to the

This transforms the nature of the language learning task: for the learner already has rich, structured, semantic contents available to communicate,

including possibly questions, goals, and plans,

This contrasts with theories of language learning that assume the child has to learn both how to mean and what to mean at the same time as learning how to communicate meanings.

It also allows learning an external language to be a collaborative creative problem-solving task in which the children are normally in a minority and “politically weak”.

Contrast the Nicaraguan deaf children

More on inner languages

GLs are not restricted to linear strings of symbols or to Fregean languages using a syntactic form composed entirely of applications of functions to arguments.

(Sloman 1971) suggested long ago that analogical representations using other modes of composition are sometimes useful for representing and reasoning about spatial configurations among other things.

Analogical representations, including diagrams and maps are capable of supporting structural variability and (context sensitive) compositional semantics since parts of diagrams can be interchanged, new components added, etc.,

But they don't have to be **isomorphic** with what they represent, as should be obvious from the fact that 2-D pictures can represent 3-D objects (e.g. the Necker cube).

The relationship is more subtle and complex than isomorphism, and can be highly context sensitive.

Internal GLs may use analogical representations not yet known to science.

The use of such representations externally (e.g. on paper, in 3-D models) usually has to be learnt or developed – the representations only work for people, animals, or machines that have suitable information-processing mechanisms.

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

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?

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. (Compare Gibson and Pick, 2000)
- 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 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).

Learning About Causation

Discussion notes and presentations on the CoSy web site

Claim: the ability to manipulate (possibly analogical) representations of spatial structures can be the basis of a kind of causal competence that enables a reasoner to understand why a certain event or process **must** have certain effects.

This uses a Kantian conception of causation that involves more than mere reliable correlation: there is a geometrical necessity in the relation between cause and effect.

It is not easy to determine what forms of representation and inference are used in animals (or children) that cannot talk.

Some of the problems of investigating causal understanding in non-human animals, were discussed by Jackie Chappell in her WONAC presentation:

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

Also Gibson and Pick (2000).

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 novel ways.
7. Some high performance competences may not scale out.
8. Many novel combinations of human competences arise in 3-D spatial processes: Spatial coincidence and proximity (with and without temporal proximity) supports novel combinations of initially unrelated structures and processes:
some animals acquire orthogonal recombinable competences
9. We do not yet understand what **forms of representation** of 3-D structures and processes (and related affordances) humans and other animals use.
10. A human without a normal body nevertheless inherits mechanisms that evolved in normal human bodies:
human cognitive development, requires embodied ancestors more than it requires embodiment.

Conclusion

We should not assume that the problem is simply a collection of inadequate tools, and immediately start trying to develop new tools, which has happened repeatedly in the history of AI.

The real gap in our knowledge is understanding of the **problems**, or, in engineering terms, the different sets of **requirements** that need to be satisfied by working systems of different sorts.

That includes, but is not restricted to, understanding the problems of learning and acting in a richly structured, changing 3-D environment.

Without doing the requirements analysis, building new tools, even biologically inspired tools, can lead us up blind alleys.

THANK YOU!

For the importance of virtual machines and supervenience see

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

Why robot designers need to be philosophers (and vice versa)

For ideas about how machines or animals can use symbols to refer to unobservable (theoretical) entities see

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

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

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

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

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

For a challenge to learning theorists: how to do dimensionality **expansion** to invent 3-D explanations for 2-D experiences (rotating necker cube):

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/nature-nurture-cube.html>

Requirements for going beyond sensorimotor contingencies to representing what's out there
(Learning to see a set of moving lines as a rotating cube.)

Invited presentation to machine consciousness symposium AAAI 2007 fall symposium:

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

Additional papers and presentations, including presentations on causation

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

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

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

[Especially papers and presentations by Chappell and Sloman](#)