

Language and Cognition Seminar

School of Psychology, University of Birmingham, 5 Nov 2004

Varieties of Meaning

Aaron Sloman

School of Computer Science

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

These slides will later be added to

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

The talk was originally entitled 'Varieties of meaning in perceptual processes' but I did not manage to get to the perceptual processes part, being developed here:

<http://www.cs.bham.ac.uk/research/cogaff/sloman-vis-affordances.pdf>

Produced and presented on a Linux-only machine

Overview – topics planned but not all covered

- **Rapid survey of theories of meaning,**
 - naive innatism/nativism (Plato)
 - naive concept empiricism (refuted by Kant)
 - structure-based theories in 20th century philosophy of science,
 - less naive innatism/nativism (boot-strapping of structures)
 - naive symbol-grounding theory (concept empiricism recently reincarnated) and the role of embodiment
 - symbol grounding vs symbol attachment
 - meanings in different types of organisms and machines (from protozoa to professors...)

Overview – topics planned but not all covered

- **Rapid survey of theories of meaning,**
 - naive innatism/nativism (Plato)
 - naive concept empiricism (refuted by Kant)
 - structure-based theories in 20th century philosophy of science,
 - less naive innatism/nativism (boot-strapping of structures)
 - naive symbol-grounding theory (concept empiricism recently reincarnated) and the role of embodiment
 - symbol grounding vs symbol attachment
 - meanings in different types of organisms and machines (from protozoa to professors...)
- **Meaning in dynamical systems**
 - The implicit/explicit distinction (not the same as unconscious/conscious)

Overview – topics planned but not all covered

- **Rapid survey of theories of meaning,**
 - naive innatism/nativism (Plato)
 - naive concept empiricism (refuted by Kant)
 - structure-based theories in 20th century philosophy of science,
 - less naive innatism/nativism (boot-strapping of structures)
 - naive symbol-grounding theory (concept empiricism recently reincarnated) and the role of embodiment
 - symbol grounding vs symbol attachment
 - meanings in different types of organisms and machines (from protozoa to professors...)
- **Meaning in dynamical systems**
 - The implicit/explicit distinction (not the same as unconscious/conscious)
- **Various intermediate cases**

Overview – topics planned but not all covered

- **Rapid survey of theories of meaning,**
 - naive innatism/nativism (Plato)
 - naive concept empiricism (refuted by Kant)
 - structure-based theories in 20th century philosophy of science,
 - less naive innatism/nativism (boot-strapping of structures)
 - naive symbol-grounding theory (concept empiricism recently reincarnated) and the role of embodiment
 - symbol grounding vs symbol attachment
 - meanings in different types of organisms and machines (from protozoa to professors...)
- **Meaning in dynamical systems**
 - The implicit/explicit distinction (not the same as unconscious/conscious)
- **Various intermediate cases**
- **How different kinds of meaning are used in different parts of a complex architecture**
 - In particular: varieties of meaning in perception and action subsystems

Overview – topics planned but not all covered

- **Rapid survey of theories of meaning,**
 - naive innatism/nativism (Plato)
 - naive concept empiricism (refuted by Kant)
 - structure-based theories in 20th century philosophy of science,
 - less naive innatism/nativism (boot-strapping of structures)
 - naive symbol-grounding theory (concept empiricism recently reincarnated) and the role of embodiment
 - symbol grounding vs symbol attachment
 - meanings in different types of organisms and machines (from protozoa to professors...)
- **Meaning in dynamical systems**
 - The implicit/explicit distinction (not the same as unconscious/conscious)
- **Various intermediate cases**
- **How different kinds of meaning are used in different parts of a complex architecture**
 - In particular: varieties of meaning in perception and action subsystems
- **Implications for psychological research and explanatory model-building (if there's time)**

Background

Various things lie behind this talk:

- Hearing people in psychology, AI, animal cognition, philosophy, asking **how do, how should, how can, people (children, adults, ...), other animals, robots, other machines, represent such and such, or how does the visual system represent X or Y?**

Can we unpack some of the presuppositions underlying such uses of the notion of 'representation'?

For philosophers this is the problem of explaining how intentionality can exist.

- The recent rise of so-called symbol-grounding theory, and related reincarnations of 'concept empiricism'. (Gärdenfors on 'Conceptual Spaces' ?)
- Much recent discussion of the role of embodiment in cognition, often linked to emphasis on dynamical systems.

(extreme – crazy – version: design the body right and cognition isn't needed!)

For an interesting collection of papers, see special issue on 'Situated and Embodied Cognition,' Editor Tom Ziemke, *Cognitive Systems Research*, 3,3 Dec 2002, <http://www.elsevier.com/locate/cogsys>

- Growing recognition of the importance of affordances: what are they and how are they represented in perceivers of affordances?

This requires representations of things that do not exist but could exist – a special kind of meaning: e.g. representing possible future actions or action sequences.

Why there's a problem

I have the impression that the discussions and debates suffer from an unacknowledged theoretical vacuum – there are many partial, implicit, theories taken for granted as the obvious, whole truth.

(E.g. meanings must come from experience? Where else could they come from? Wrong!)

Much relevant work has already been done, e.g. by philosophers of science (described below), but ignored by many in the field.

We need an overview, a better understanding of what kinds of theories of meaning are possible, leading to a more principled general framework for the detailed research questions that involve meaning.

The following assumption is a good basis for making progress:

All living things process information, in the sense of 'information' that involves 'meaning' or 'reference', rather than the Shannon sense.

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

We need to understand the variety of types of information and information processing in organisms and machines. I'll start with typical linguistically-expressed meanings (explicit information) because so much has already been done in this area, and then move onto implicit information.

A very old philosophical problem

How can thoughts, symbols, words, pictures, brain states, or anything else refer to something?

- The oldest answer (in 'western' philosophy) is perhaps Plato's notion of 'memories' of another world, the world of ideal forms.

This naive innatism/nativism begs the question of how reference worked in the previous life and how it works in this life – how can memory of a previous life help?

- Another old answer is **concept empiricism**:

This claims that all meaning comes from experience of instances of concepts – each concept is formed by abstraction from experiences. (how??)

E.g. the 'British Empiricists': Locke, Berkeley, Hume, and many others.

Most people not trained in philosophy regard concept empiricism as obviously true because it 'seems obvious' and they cannot imagine any alternative to it.

- Kant (circa 1781) demolished concept empiricism.

You can't get concepts from experiences alone, since having experiences requires the use of concepts (e.g. concepts of spatial and temporal relations, and perhaps the concept of cause).

So he suggested some concepts must be innate. Chomsky, thinking about language learning, reached a similar conclusion in the 1960s, though details of his theory have changed.

Is there any version of innatism/nativism that can be shown to have explanatory power and that can be tested in a robot?

20th Century Developments

Philosophers of many kinds played with variants on concept empiricism, or when those failed, various kinds of relativism, e.g.:

- Logical positivism/Verificationism (The 'Vienna circle')
- Operationalism (P.W. Bridgman – Winner Physics Nobel Prize 1946)
- Wittgenstein: **Don't ask for the meaning ask for the use.**
- Kuhn's and Feyerabend's versions of relativism
- Popper:

Meaning is not worth discussing, only truth/falsity and relevance to the goals of science.

None of these philosophical approaches explained how a meaning-user could actually be designed and built.

They assumed such things existed (e.g. adult humans), then offered criteria for identifying 'meaningful' (or in Popper's case, 'scientific') propositions and theories.

The authors were not interested in how other animals might be meaning-users, without using language, or how prelinguistic human children might develop into language users, or how machines might have intentionality.

But a very important idea grew out of the difficulty of explaining how scientific theories referring to unobservable entities (e.g. sub-atomic forces, genes, economic inflation, etc.) can be meaningful.

Structure-based theories of meaning

Some thinkers noticed that in science and mathematics much meaning comes from **structure**, e.g. the structure of a theory implicitly gives **some** meaning to the undefined/primitive notions used in the theory.

In this sense a theory includes

- Some form of representation, or collection of formalisms, in which assertions, questions, predictions, explanations can be expressed, including rival theories.
- Some theory-independent means of manipulating representations in order to make inferences, check for consistency, compare questions and answers, etc.: **truth-preserving or denotation-preserving transformations.**
- Labels for 'primitive' concepts, used in the theory without explicit definitions.
- 'Axioms' (assumed, or conjectured truths) expressed using the formalisms, which constitute the claims of the theory, from which others can be derived.

A question for later:

How many varieties of formalisms and types of manipulation are there, and how do they differ in what they are good for?

Think of logic, algebra, grammars, computer programs, diagrams, maps, genomes, neural nets, wave forms...

<http://www.cs.bham.ac.uk/research/cogaff/sloman-analogical-1971/>

How can structure determine meaning or reference?

The key idea is that structures map on to other structures in various mathematically specifiable ways, and in that sense one structure, M , e.g. a portion of the physical or social world, or some mathematical entity, can be a **model** for another structure, S , a collection of sentences, equations, diagrams, etc.

Tarski made this old mathematical idea precise when S is set of predicate logic sentences

The way in which the structure S constrains the set of possible models M gives S a kind of meaning.

A user of S , e.g. an animal or machine containing S , may exploit those structural relations in its perception, action and reasoning mechanisms.

Philosophers found that obvious but did not develop the idea in detail.

Examples follow.

Adding meaning by enriching structure

- Adding an (independent) axiom, or constraint, to an existing theory will reduce the set of possible things in the universe that can be models of the theory.
- But that can never **guarantee** that the theory refers definitely to a unique bit of the world, for there could always, in principle, be a sufficiently similar duplicate bit somewhere, which is also a model (**the 'twin earth' argument**).
- So Strawson proposed (in *Individuals* 1959) that causal connections are needed to pin down reference.
- However (as John McCarthy pointed out in conversation) as a theory gets more complex the probability of duplicate models is reduced, and may be so low that it can be ignored.
- In any case, we don't need what philosophers (e.g. Strawson and philosophers who regard rationality as a requirement for intentionality) seek, namely a way of **guaranteeing** uniqueness of reference: in real life it suffices if there just **happens** to be only one thing, without any guarantees.
Success without guarantees is still success!

So in principle, structure can suffice in practice to determine reference.

See my papers on this in IJCAI 1985 and ECAI 1986, both here

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

However, there are other ways of constraining reference, namely ...

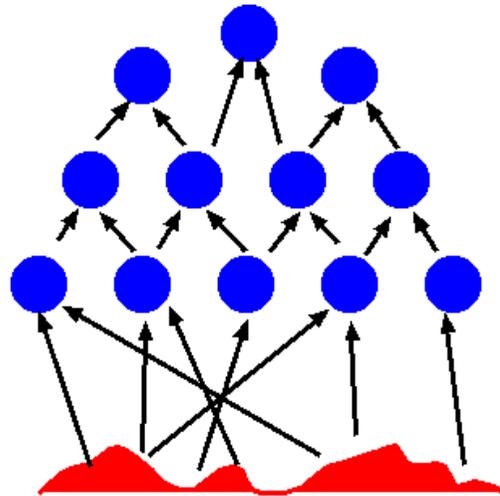
Varieties of attachments between theory and world

Besides enriching the structure (adding 'axioms') there are various ways of making what a theory refers to in our world more definite:

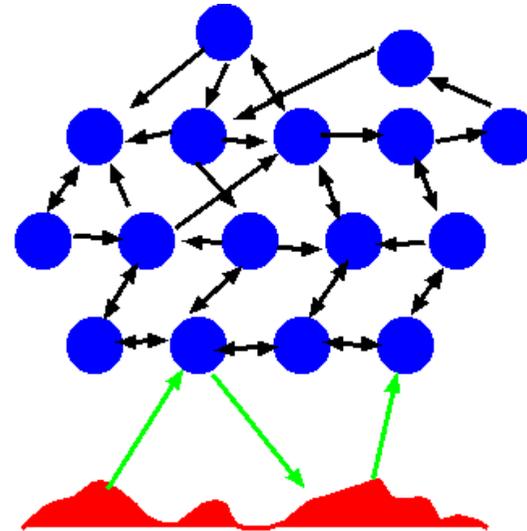
- 'Bridging rules' may **explicitly** define a subset of the primitive theoretical terms in terms of other already meaningful terms
(E.g. operational definitions: 'electric current' defined as what is observed by building such and such a device, and using it in such and such a way in electric circuits. But means of measuring current can be criticised and replaced, and after replacing old devices with new we still say we are measuring current, for most of the meaning of 'current' comes from the theory containing it – not from bridging rules, measuring devices/procedures, or experiences.)
- Carnap's notion of meaning postulates, which **partially** define new terms by expressing their relations to old ones.
Something is soluble in water **IMPLIES** if were X to be done Y would happen – this sugar is soluble because **IF** it were placed in water and stirred, then.....
(heavy reliance on counterfactual conditionals – something untested could be soluble)
- Some of these ideas lead to the Quine/Duhem notion of theories touching reality only at their fringes, with a core that's immune to direct verification or falsification by observation or experiment.
Mismatches at the fringes of a theory perturb the central tenets of the theory to varying degrees – mostly not at all: we repair the fringe normally. But accumulating problems may require a massive change (e.g. relativity theory, quantum theory, darwinian theory,)

We'll see how developments in a child or robot could mirror progress in science

Symbol Grounding vs Symbol Attachment



Symbol Grounding



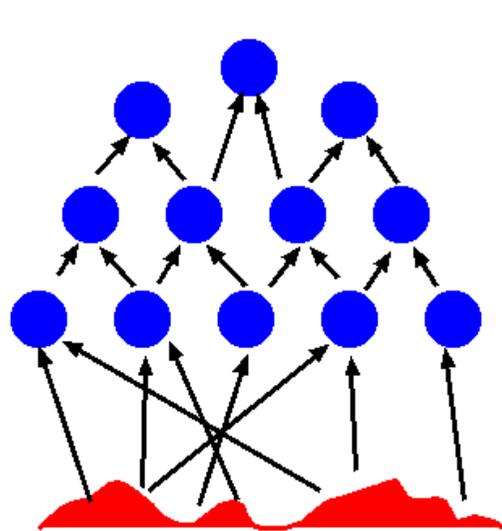
Symbol Attachment

The idea of **symbol-grounding** (concept empiricism) is that all meaning is derived from ('flows upward from') sensory processes (experiences), whereas the idea of **symbol attachment** is that a great deal of meaning comes from the structure of a theory and internal constraints, limiting possible models. Bridging rules (links with perception and action), attach the structure to the environment and thereby help to reduce or remove residual indeterminacy of reference and make the theory applicable.

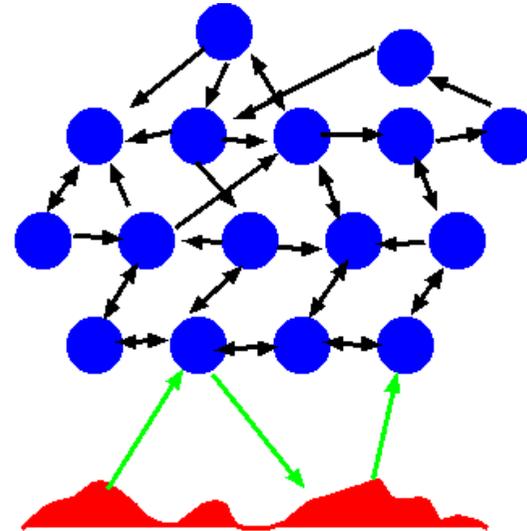
See <http://www.cs.bham.ac.uk/research/cogaff/talks/#grounding>

It's more complex than that: we have many mutually supporting theories. Moreover, because a theory uses a reasoning system, new attachments (e.g. types of experiment or application) can be derived from old ones: the source of explanatory and predictive power of theories.

Change of terminology: Attachment → Tethering



Symbol Grounding



Symbol Tethering

Note added 23 Aug 2005

**I have been persuaded by Jackie Chappell
to replace the phrase 'Symbol Attachment'
with 'Symbol Tethering'.**

Why a child, or a chimp, needs deep theories

Perceivers like us need some understanding of how shapes, colours, weights and other properties of things change as various things are done in the world, and how our experience of them should change while we move or act and the context (e.g. lighting) changes, though the objects persist unchanged.

Kant: we have to assume there is a reality whose properties persist independently of our changing experiences of them – as the hidden properties of atoms, sub-atomic particles, force-fields, genes, niches, evolutionary pressures, economic and social forces, etc. persist independently of our changing experiences of them.

Something like this seems to be a requirement for understanding some of the most mundane aspects of our world and some of the deepest.

It is often assumed that all such knowledge of the environment must come from experience, by various kinds of induction or abduction, or statistical inference.

That assumption could be false if evolution produces structures that constrain what the individuals of certain species are capable of learning, so that not all details of the theory of the environment have to be learnt empirically (as in precocial species), or new theories, not derived from experience, are invented.

Note: whether what is learnt or innately determined is **implicit** (only embodied in the design of the system) or **explicit** (represented in some formalism that the system can manipulate, e.g. in making inferences) is a separate issue.

In microbes and insects it is probably all implicit. Some mammals and birds may be different.

Let's revive an old proposal

A theory of meaning that is good for scientific concepts referring to deep, unobservable entities used in theories with rich explanatory power, may also be a good theory for the kinds of meanings often claimed to be derived from experience, e.g. concepts like 'red', 'rough', 'rigid', 'heavy', 'smooth', 'sweet',

All these familiar concepts are embedded in very old implicitly used explanatory and predictive theories about the environment.

Perhaps more of their meaning comes from their structural roles and relations in those theories than from the experiences associated with them.

Perhaps 'red' is more like 'quark' than people think.

Naive innatism almost fits precocial species

If genes can produce at birth appropriate structures (theories), and mechanisms for using them, then, insofar as those structures allow some things in the environment as models and others not, the structures may determine much of what the concepts used connote.

So innately determined structures (used via links to sensors and motors) may provide all or almost all the meaning required by members of a precocial species. Indeterminism and ambiguity of meaning may be removed or reduced by (implicit) rules linking innate structures to perceptual and action mechanisms.

These help to pin down the interpretation of the theory so as to make it refer to the immediate environment, rather than some isomorphic 'twin earth'.

- This works dramatically for precocial species, e.g. deer that can run with the herd within minutes of being born.
- They don't have time to develop all the conceptual apparatus on the basis of their individual experience.
- Innate structures and manipulation mechanisms can enable a novel configuration of terrain (e.g. a rock ahead) to be understood in such a way as to produce appropriate action – e.g. making a detour or climbing over.

What about altricial species?

Beyond naive innatism: altricial species

For altricial species – e.g. humans, primates, hunting mammals, nest-building and hunting birds, the genes do not provide all the generic information used in adult life (for reasons mentioned below).

Instead there seems to be a genetically produced ‘bootstrapping program’ that allows the structures to be built during early interaction with the environment, even while the brain is still growing. This may explain why certain kinds of learning have to occur at critical ages (approximately).

That means the later stages are a product of both nature and nurture, and it is possible for nature to impose strong constraints on what nurture can generate: millions of years of evolution should not be wasted.

The implication of this would be that the theoretical structure determining the meanings used by an adult of such a species could have many abstract (high level) features determined genetically and common to all members of the species, whereas concrete details are determined by the individual’s environment (and culture).

(Compare Chomsky and others on language universals.)

NOTE: Genes may be unable to provide full information because

- there is too much to encode explicitly in the genome – post-natal inference is more economical.
- there never was a time in the past when certain information could have been acquired
- a generic, creative, mechanism is more versatile in a changing world than fixed information.

Some intermediate cases

Terminology is often confused and inconsistent in discussions of these matters.

- A system able to construct and manipulate explicit meaningful structures of unbounded complexity, to construct plans, predictions, hypotheses, questions, explanations, using compositional semantics, is often described as having **deliberative** mechanisms, because it is capable of hypothetical reasoning.
- But there are many intermediate cases in the evolutionary history, to which different people apply different labels.
- E.g. some people use the word 'deliberative' to describe a reactive system in which two or more action tendencies (e.g. flee or fight) can be simultaneously activated, with a mechanism to make sure only one of them wins out. I prefer to call these **proto-deliberative** mechanisms, if they don't include the full range of capabilities listed on previous slides.

We need to investigate many intermediate cases including:

- **more or less explicit representations of goals or needs**
- **more or less explicit representations of internal context**
- **more or less creativity in the organism's mechanisms for combining information of different kinds.**

Microbes and insects are not like chimps

Does a microbe or an insect need a deep theory of the environment?

Not if a set of pre-compiled reactive behaviours suffices to enable it to survive and procreate....

- Many purely reactive organisms are able to achieve great things **in the right environment** using only implicit information in patterns of activation in a fixed architecture implementing a large number of pre-compiled condition-action rules, e.g. in a web-building spider, or a cathedral-building termite.
- Organisms that need to find solutions to new problems by thinking ahead, need the kind of understanding that makes use of explicit manipulable representations that go beyond the information implicit in a fixed collection of condition-action rules,
- Greater generality and flexibility comes from use of a structured theory with transformations or inference mechanisms for applying the theory to new contexts.

However, note the costs of both the working memory mechanisms, with a garbage-collectable store of temporary structures, and the longer term content-addressable associative memories required for planning or predicting, especially reasoning several branching steps ahead.

Proto-deliberative mechanisms (e.g. competitive networks) manage only one step planning.

Implicit meaning/information in reactive systems

- In many organisms and machines the structure determining meaning is not a formal theory, but **the architecture of the machine** and its mode of interaction with the environment, using sensors, effectors and various kinds of internal control states and processes.
- E.g. an array of sensors feeding signals into a neural net may be capable of activating different patterns of control signals to the effectors.
- The whole thing is a constantly changing **dynamical system**.
- All the information about the environment is represented only in the current patterns of activation within the network of control signals (including excitation and inhibition) — information is there only insofar as it is causally **active**.
- Everything is **transient** – when the causes of activation die, the activation dies and with it the information expressed.
- Information may be **composed** in different ways as different patterns of input activation spread through the system, possibly triggering winner-takes all subnets.
- But the modes of composition of information in a purely reactive system using only implicit information are limited to superposition and blending, since those are the ways of composing activation patterns.
- (Sequencing becomes possible if an **explicit** changeable state representation (e.g. a counter) is sensed in conditions and altered in actions.)

Implicit particular and general meanings

Some purely reactive systems with only implicit episodic information can learn. This uses 'architectural changes' such as varying connection strengths to encode general information.

- **If there is learning/adaptation, e.g. by changing connections or connection strengths, general, reusable, information is implicit in the network.**
- **It may implicitly express generalisations like 'As are followed by Bs', or action rules like 'In context C if A is sensed do X otherwise do Y'**
- **Such general information manifests itself only through its causal role in processing inputs and generating or selecting outputs.**
- **This contrasts with implicit information about particulars (episodic information) expressed only in activation states stimulated by sensing those particulars**

We can therefore see how implicit information in insects and many other reactive organisms and machines can be expressed in a mixture of activation states and weak and strong activation linkages all causally embedded in an environment (including some internal states of the organism) capable of being sensed and acted on.

Those organisms manage well without explicit information.

Steps toward explicit meaning

Biological evolution 'discovered' the usefulness of an additional form of representation: something that can endure beyond the initial activation, and can be re-used and recombined later in different ways.

- Important new powers came from going beyond physical states of activation implicitly encoding information to the use of stored enduring explicit structures (including structures in virtual machines) that can be created, erased, copied, manipulated and combined in various ways with other structures, using compositional semantics, (with a major cost in 'hardware', not discussed here).
- E.g. it became possible to construct
 - descriptions of desirable future states, or goals: (e.g. that berry in my mouth, these children in that cave, that predator chased away, food grown in the future)
 - hypothetical explanatory histories for observed states: (these marks could have been caused by a hungry predator walking here)
 - hypothetical descriptions of unobserved parts of the environment (something dangerous, or something edible, outside the door of the cave),
 - predictions of future states (that predator will get to the children soon)
 - sequences and other combinations of possible future actions producing a goal represented in advance (a sequence of actions needed to get that food in my mouth, or to get me between children and predator).
 - Hypothetical reasoning: if I do X then Y will not happen, etc.

I.e. evolution 'discovered' the importance, for at least some organisms, of moving beyond mere dynamical systems long ago.

‘Precocial’ does not imply ‘reactive’

- Organisms that manipulate only implicit information expressed in patterns of activation of physical structures are limited in how they can vary what they know or do, for they have no means of developing complex new behaviours, though they may, through patterns of positive and negative reinforcement, change their ‘contingencies’, a useful form of learning in environments which change in simple ways.
- Thus most of their behavioural capabilities must be predetermined in their physical structure and neural wiring. I.e. their information handling capabilities will be mostly genetically determined.
- In many cases they will be ready to use those capabilities unaided at birth.
- This is desirable if their parents lack the ‘intelligence’ or the physical power to look after them.
- Thus purely reactive organisms (manipulating only implicit information and lacking deliberative capabilities, apart from proto-deliberative ones) are likely to be precocial.
- **But it is possible for precocial organisms to be provided by their genes with mechanisms for some planning, prediction and creative control of behaviour on the basis of explicit information mechanisms.**
So some not purely reactive species can be precocial.
Grazing mammals, e.g. deer, may be examples, if they can think more than one step ahead.

Organisms may have old and new mechanisms

- More complex mechanisms need not be entirely reactive, nor entirely deliberative.
- Many bodily functions including things like posture control, temperature control, digestion, protective reflexes (e.g. blinking, wincing) use reactive mechanisms that evolved aeons ago and function well, with only **implicit** information.
- That is consistent with also having mechanisms that acquire, manipulate, store, and use **explicit** information structures and are capable of hypothetical reasoning of various kinds, and varying degrees of sophistication.
- So the information processing architecture of an organism, or a machine, can include a mixture of deliberative and reactive mechanisms of many kinds.
- Moreover, those mechanisms may have to share sensors and motors, e.g. sharing visual or auditory input between ancient reflex mechanisms and sophisticated language understanding capabilities.
- It is also possible for a particular skill that starts reactive to become understood by the user of the skill and subject to deliberative control later on. Conversely, repetitive use of a skill that is initially only under deliberative control because it uses the result of new planning or new learnt instructions from another agent, may train a reactive sub-system to take over control and perform the actions faster and more fluently, leaving cognitive mechanisms for other tasks.

Beyond insects: is it all nurture?

- Evolution moved beyond insects and other purely reactive systems.
- But that does not imply that it provided humans with **totally general** knowledge-free learning mechanisms, able to learn absolutely anything by induction from experiences.
 - **Could we bring up infants talking in predicate calculus?**
- Both the physical forms of sensors and motors and the types of virtual machines that specific brain structures support can constrain what is capable of being learnt

The genes may have their say indirectly, imposing powerful constraints on what sorts of forms of representation, concepts, theories are capable of being learnt.

(Millions of years of evolution are not wasted.)

Beyond Insect-like meanings

Evolution produced deliberative and meta-semantic capabilities

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

- Deliberative mechanisms need non-transient, explicit, forms of representation, which can be treated as objects to be manipulated, e.g. in generating hypotheses, predictions, plans, etc.
- Metamanagement needs meta-semantic representations of systems that themselves represent (process information)

Varieties of meaning in a complex architecture

We need to develop theories of how different kinds of meaning are used in different parts of a complex architecture.

In particular: varieties of explicit and implicit meaning in perception and action subsystems, each of which may include different layers performing different tasks at different levels of abstraction.

Meanings of different kinds (including more or less explicit, structured, varied) can occur within motivational and other affective mechanisms.

There are plans to develop these ideas within the CoSy project:

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

Where I've got to

This is work in progress with much still to be done.

a changing snapshot of the ideas is in a draft, growing, disorganised paper here:

<http://www.cs.bham.ac.uk/research/cogaff/sloman-vis-affordances.pdf>

Offers of help always welcome!

This slide presentation presents the material from a different perspective.

Yet another perspective is this set of slides attacking the notion of 'symbol-grounding'

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

See also the tutorial slides on integration presented at the EC Cognitive Systems Kickoff Conference, in Bled, October 2004

<http://www.cs.bham.ac.uk/research/projects/cosy/bled-o4.pdf>

To be continued – perhaps