



Evolved Cognition and Artificial Cognition

Some genetic/epigenetic trade-offs for organisms and robots

Or: Notes on how to do AI-Inspired Biology (AIIB)

(INCOMPLETE DRAFT LIKELY TO BE REORGANISED)

Aaron Sloman

(Philosopher/Cognitive Scientist/Programmer)

School of Computer Science, University of Birmingham.

Last updated: 31 Jan 2011

Installed: 30 Jan 2011

This file is <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evolved-cognition.html>

Also available as PDF (created using html2ps and ps2pdf).

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evolved-cognition.pdf>

The PDF version may sometimes be a bit out of date.

Introduction: the "tabula rasa" theory

It has often been assumed (even by Alan Turing) that it is possible to design a machine that will acquire human-like intelligence if it initially has

- (a) plenty of information storage capacity,
- (b) a very powerful general-purpose learning mechanism,
- (c) a rich environment in which to learn, and
- (d) a teacher to guide the learning.

The implicit claim is that such a machine can learn at speeds comparable to learning in humans, and will not require evolutionary time-scales, despite having no specific knowledge about the environment initially, nor even any concepts (an ontology) specific to the environment. This is closely related to the ancient "tabula rasa" theory of knowledge acquisition, summarised here:

http://en.wikipedia.org/wiki/Tabula_rasa

For some who do not study human infants closely it seems plausible that that is how human learning works, since neonates appear to be almost completely incompetent and ignorant, to start with.

A clue that something is wrong with this "tabula rasa plus learning machine" view (TRLM) comes from the observation that so many species produce offspring that either have to fend for themselves completely, or at least have to keep up with adults.

For instance, without being taught:

- birds extricate themselves from the egg when ready to hatch,
- young ducklings identify a nearby moving object, with some adult-duck-like qualities, to imprint on and follow
- many invertebrates live two entirely different lives, e.g. one larval and one as flying insect, and do not have to **learn** how to live either type of life, nor how to transform themselves from one form to another
- the young of many grazing mammals are able, within minutes or hours, to get up, walk to the mother's nipple, suck and, in some species, soon after that run with the herd, e.g. to escape predators: displaying visual and other competences no current robot comes even close to

If so much competence assembled over evolutionary time-scales can be transmitted (presumably mostly encoded in genomes, or genomes plus their epigenetic environment, e.g. features of the mother's womb) to so many species, is it possible that people who think humans are born incompetent and ignorant we are missing something very important?

Some alternative hypotheses concerning innate knowledge and competences

To contrast with the TRLM theory, I shall outline some alternative hypotheses regarding starting points for various kinds of learning about the world in diverse animals and, by implication, future intelligent robots.

The hypotheses differ according to what they assume the products of biological evolution available from birth or hatching are, and what they assume about how we can find out. Although this is somewhat oversimplified, I shall group the main answers into the following categories (labels are provisional):

- **The empirical nativist thesis (ENT)**
(Currently manifested in a lot of work in developmental psychology.)
- **The rationalist nativist thesis (RNT)**
(One of the old ideas in philosophy, going back to Plato and probably earlier.)
- **Empirical meta-knowledge nativism (EMN)**
(Closely related to Chomsky's ideas about linguistic universals).
- **Design-based nativist meta-knowledge (DNM)**
(Presupposed by what McCarthy calls "The Designer Stance".)

The empirical nativist thesis (ENT)

During the last two decades an empirical hypothesis strongly opposed to the TRLM theory has emerged among developmental psychologists, namely that there is a considerable amount of **innate** knowledge in humans though its presence is not immediately obvious. So an "empiricist nativist" industry, led by Elizabeth Spelke, among others, has been probing infants at various stages in their early development for signs of "core concepts", innate competences, etc.

The rationalist nativist thesis (RNT)

The ENT echoes a much older view among philosophers, including Plato and Kant, that humans could not have experience and acquire knowledge empirically if they did not have certain kinds of concepts and knowledge to start with.

Unlike the more recent empirical nativists, rationalist nativist philosophers think that it is possible to work out from "first principles" what the innate knowledge must be, e.g. starting from conceptual analysis, without doing empirical research to find out what the innate knowledge actually is.

Empirical meta-knowledge nativism (EMN)

A more sophisticated theory-driven type of research, often inspired by some of Noam Chomsky's ideas, assumes not that the neonate's knowledge displayed in successful behaviour is innate, but that some highly schematic meta-knowledge, about a class of environments, is innate, which is instantiated to specific knowledge and competences through interaction with the environment after birth or hatching. The knowledge the child actually needs in its life is derived by transforming the schematic forms to substantive instances by acquiring parameters from the environment to produce the required instantiations. (Chomsky described himself as a Cartesian, whereas his views seem to me to be closer to Kant's, despite the difference highlighted here.)

This is a generalisation of Chomsky's thesis that individual human languages differ in detail but all are instantiations of a universal language schema specific to humans, and specified in their genome.

The diversity of human languages, all presumably products of some shared innate human-language learning competence (a set of hypothesised "language universals") is sometimes taken to illustrate this, though not all researchers agree with Chomsky that the genetic component that gives rise to language is language-specific. (The search for language universals has not been very successful.)

Generalising Chomsky's idea to include visual competences, motor competences, learning competences, reasoning competences, social competences leaves open the possibility that human language learning does not have its own genetic specification, but is largely based on a combination of other more general generic schemata.

This generalised EMN can inspire empirical research on humans and other animals to find out just what the innate, parametrisable, meta-knowledge is, as opposed to research aimed at finding out what the innate concepts and knowledge of a neonate are.

On this sort of view the innate knowledge, though abstract and missing detailed parameters is not totally general: it is suited to development by certain forms of interaction only in a restricted class of environments, a class that generalises features of the environments in which the species evolved. In that case the innate knowledge would share many characteristics with other species that share some of that evolutionary history. So, for example, the learning mechanism might be incapable of driving much learning in a baby whose sensory input from the environment is restricted to something like a 2-D TV display of the view from a mobile camera in a field, and whose only motor signals are ones that can alter the position, orientation, and motion of the camera.

The precise details of the EMN mechanisms and innate meta-knowledge would have to be discovered empirically by revising and extending the techniques of Spelke and others so as to investigate not the specific innate concepts and knowledge, but the innately specified patterns of

development of concepts and knowledge.

Design-based nativist meta-knowledge (DNM)

A somewhat different approach (apparently triggered by reading the work of Spelke), was suggested by John McCarthy, in "The well-designed child" (originally written in 1996, but published with minor revisions in 2008). I shall try to articulate some features of his DNM approach that I don't think he made fully explicit -- though there is a danger that I am hallucinating some of my own ideas not what he wrote or intended.

Like those who assume EMN, he starts from the observation that products of biological evolution must have been shaped to a considerable extent by features of the environments in which they evolved -- which may have had some constant and some changing characteristics.

So, instead of (a) trying to derive the nature of the innate component from totally general abstract principles (e.g. principles of rationality), or (b) trying to discover it purely empirically by observing as many as possible of the exercises of competence in very young children, McCarthy proposes (c) that we should study the environment or environments in which a species evolved, in order to try to **work out**, on the basis of general features that are common to those environments, and possibly also features of some of the early environments in which the evolutionary process occurred, the requirements for being able to do various things, including learn and reproduce, in those environments.

On the basis of those requirements we can try to design explanatory theories of mechanisms that meet those requirements and test them by building working models. In practice, of course, the discovery of requirements and the process of design and testing have to be done in parallel with much mutual influence: since some requirements only become apparent after partially working designs have failed.

The scientific investigation based on DNM therefore overlaps with the engineering activity of trying to design a machine that performs in such an environment. McCarthy used the label "the designer stance" to characterise such research: it is related to, but slightly different from what Dennett called "the design stance".

DNM differs from EMN insofar as it relies less on empirical observation of very young humans and other animals, while not discarding empirical observation, since that is needed to test the theories developed. However, a deeper test than formulating and testing predictions from a theory that is usually not very precisely specified, is building working instances of the theories to demonstrate that they can in principle do what is claimed. Typically, creative designers will be able to come up with alternative, competing, theories about innate meta-competences and empirical research will have to be added to demonstrations of working systems to test the theories. (See Sloman&Chappell IJCAI 2005, and Chappell&Sloman IJUC 2007 for some outline ideas about the forms of DNM theories.)

DNM therefore leads to a mixture of empirical research combined with creative design, such as is required for building machines that work in complex, varied, and changing environments. However some of the empirical research feeding into this work will be research into features of the environment that both influenced biological evolution and also support the developmental processes driven by evolved mechanisms.

Compared with the purely empirical approaches of ENT and EMN, and the purely philosophical arm-chair approach of RNT, using the designer stance seeking DNM theories can be a powerful source of new, testable, ideas about what evolution might have provided, especially if researchers doing that research have the benefit of real experience of designing, building, testing and debugging working systems that provide insights into unobvious features of the problems evolution may have solved. (The philosophers and psychologists usually have not had any such experience -- and it shows in their theories.)

Actually, at least one psychologist, Ulric Neisser, suggested that research on cognition does not take enough account of the environment in his 1976 book, *Cognition and Reality* (which I have not yet read!). James Gibson's ecological approach also emphasises general features of the environment that provide opportunities (and sometimes problems) for perceivers who act in that environment, but, as far as I know, did not recommend building working models. A great deal of Piaget's research on development, which is not concerned with innate competences but with developmental processes also pays attention to detailed, but generic features of the environment which pose challenges and opportunities for learners. It was only near the end of his life that he learnt about AI and recognised that it might have been a very useful adjunct to his theoretical and empirical work.

Those who are familiar with the philosophy of science of Imre Lakatos will understand my claim that the DNM is more likely than the others to produce strongly **progressive** (as opposed to **degenerating**) research programmes.

Implications of the Designer Stance

The implications of McCarthy's DNM approach are not all obvious, and carrying out the project is difficult. We can distinguish at least the following aspects, some of which go beyond what McCarthy himself wrote.

- a) Animals of different species can share an environment, including animals that interact with one another (e.g. predators and prey, or animals that have a symbiotic or competitive relationship). This can give clues as to how features of the environment pose information processing problems whose solutions can generalise across species of different forms. (E.g. birds, primates, elephants and octopuses can perceive and manipulate objects in their environment, in order to achieve goals.) Some of the animals that are very different in details of morphology and their competences share evolutionary histories, so they may still use common evolutionarily old sub-mechanisms, forms of representation, ontologies, extended in different ways by more recent components or features.

Looking at those species differences may help researchers to separate out design **requirements** from design **solutions**. Without that cross-species investigation theorists may tend to conflate the specifics of the designs they propose with the problems solved by those designs, failing to think of alternative designs.

For these reasons (and others), DNM-inspired research should consider more different species than researchers in AI and psychology normally consider.

- b) DNM encourages researchers to attempt not just to list actual features of environments, but also to specify what sorts of things **can** exist in the environment, about which animals or robots may need to acquire information, including: things that might need to be perceived, thought about, acted on, created, destroyed, used, eaten, avoided, communicated with, etc.

This is likely to include not only whole connected physical objects and their properties and relations, but also object parts, surface fragments, kinds of material, etc. and also many processes in which these things change their properties and relationships. (I have elsewhere discussed the importance of multi-strand relationships between complex objects, and multi-strand processes, when such relationships change.)

Most organisms move much of the time and are surrounded by things in motion as well as things that are static. Consequently it is likely that from the very beginnings perceptual subsystems had the function of acquiring information about **processes** occurring in the environment rather than just **objects** and **situations**, unlike much research on perception in AI and psychology, which starts from perception of static entities, with the intention (if the need is considered) of adding motion later. For an organism a static scene may just be a special case of a process, in which all changes are minimal (e.g. have zero velocities).

For some organisms, the environment includes not only inert or passively moved physical entities and processes, but also things that process and use information to select and execute actions -- i.e. **agents**. So, besides the semantic competences required for coping with the former, some animals (and future robots) will need meta-semantic competences for representing and reasoning about semantic competences and their uses in the latter -- i.e. for perceiving, thinking, reasoning about, and interacting with other agents. This may, or may not, include self-referential meta-semantic competences, e.g. the ability to be self-conscious regarding one's own experiences, thoughts, plans, decisions, etc.

(Which comes first, self or other representation, may be a chicken and egg question: very likely they developed in parallel, serving different but overlapping functions.)

One of the important features of the environment that many animals need to be able to cope with, but are not always taken into account by researchers is the fact that spatial and spatio-temporal structures and processes exist on very different scales, where the smaller scale entities are embedded within larger structures. For example, someone peeling a banana may or may not need to relate that action, and its consequences to other things in the field of view, or to other things in the immediate environment that can quickly be brought into view by small movements, or to other things in the larger environment that can only be perceived by large scale motion of the whole agent (e.g. the rest of the building). In addition to relating objects to larger containing structures it is often important to relate events and processes to larger containing temporally (or spatially and temporally) extended events and processes, for instance when relating the current situation to possible future plans, or possible explanations (causal histories).

These capabilities are relevant to perception as well as to thinking, reasoning, planning, and remembering. In particular a great deal of work on visual perception focuses on perception of objects (and how they are represented) ignoring perception of processes (including multi-strand processes) and perception of possibilities for processes and constraints on such possibilities. (Gibson, like many researchers on embodied cognition, focused only on a small subset of these, namely the affordances for the perceiver.)

Some of the information will be about transient, fast-changing states and processes, and some will be about relatively enduring structures and relationships on various scales (e.g. a particular cave, the immediate environment of the cave, the larger terrain containing that environment, etc.) and some will be about processes in which details are changing while some high level process structure persists (e.g. a carnivorous mammal chasing a deer).

Some information will be about how things actually are, some will be about what is possible (what changes can occur in any given configuration), and some will be about constraints on the possibilities (e.g. laws, necessities).

A host of especially deep and difficult problems with a long philosophical history concerns the nature of causal knowledge of various kinds. From a designer standpoint, Jackie Chappell and I have argued that animals need both Humean (correlational) forms of causal information and Kantian forms, where causal reasoning is structure based and non-probabilistic, e.g. something like mathematical reasoning in geometry.

- c) One the very deep issues concerns how metrical aspects of spatial structures and relationships should be represented. Most researchers (unlike Piaget) simply take for granted that visual mechanisms have access to the kind of global cartesian coordinate system that engineers and scientists now take for granted, forgetting that this was a relatively late development in human culture. So there are vast numbers of research publications attempting to represent both image structures and scene structures in terms of such coordinate systems, often assuming that structural complexity can be handled simply by using sets (or vectors) of numerical measurements. Often the sensors available cannot provide precise and accurate numerical values, causing researchers to develop elaborate theories using probability distributions and noise-reduction techniques. I suspect there are much better solutions, based on representing structures by structures often using abstraction to remove commitment to unavailable details.

There is very little evidence that biological evolution produced brain mechanisms that are capable of using global cartesian coordinate systems where locations, motions, orientations, sizes, etc., are represented by vectors of real numbers -- except for humans who can do it after extended education in modern schools and universities.

This raises questions about what alternatives are available. I have begun to explore alternatives involving mixtures of topological structure (containment, overlapping touching, and various kinds of discontinuity in structures and processes) augmented by networks of partial ordering (of relative size, relative distance, relative angular magnitude, relative curvature, relative speed, etc.) The details get very messy especially if the need to represent multi-strand processes is taken seriously, but I conjecture that this approach is capable of avoiding some of the fragility and other problems caused by imprecision and noise in sensory mechanisms. For example, imprecise information about boundaries of a pair of objects may not matter when the question is whether one of the boundaries encloses the other object. (Much work is still needed to turn this hunch into a demonstrable functioning design.)

I suspect that pursuing that idea can lead to considerable demoting of the role of probabilistic inference in perceiving and acting in the environments encountered by many animals. That's not to deny that metrical precision and probabilistic mechanisms are never important. Metrical precision is required for throwing things at small targets, for grasping small objects quickly, and for a cat jumping from the ground to land on the top of a narrow wall. However, in many other cases the need for metrical precision in perception and motor control can be eliminated by visual and other servoing (as anyone who has undressed and got into bed in the dark will know).

- d) Another deep issue concerns kinds of stuff -- varieties of material of which things can be composed. A young child encounters many different kinds of stuff including some of the different materials of which it is composed, different parts of bodies of others (mother especially), materials constituting clothing, toys, food, furniture, nappies, towels, tissues, cotton wool, sponges, water, then later things like mud, sand, plasticine, elastic bands, paper, crayons, and various parts of plants and other animals.

At present I don't think anyone has good ideas about how those different kinds of stuff and their properties (many of them dispositional properties) are represented in the minds of humans (of any age) and other animals or how they should be represented in future robots.

Researchers may be tempted to assume that the methods of simulating physical processes in computer games or in rendering simulated scenes and processes are relevant. But those forms of representation are not typically chosen to enable an active agent to perceive, think about, act in and learn from the environment: their functions are much more restrictive. In general, forms of representation that are useful for generating synthetic movies are not necessarily useful for thinking about possibilities for action and reasoning about their consequences. Moreover, even if the simulation tools are used they will produce results at the wrong level of detail. When thinking about how I might travel to a conference in another country videos of all the detailed steps in the processes of travel would be an enormous waste.

e) Virtual machinery and physical machinery:

People who are not familiar with, or have not understood, advances in tools and methods for representing information and controlling processes over the last 60 years or more often assume that all the information-bearing structures must be physical: and seek evidence for them from neuroscience. This leaves out the possibility of increasingly abstract and flexible virtual machinery containing more powerful forms of representation implemented in, but quite different from, physical structures and processes (e.g. chemical and neural structures and processes in brains). That can include both application virtual machines, that have some specific information processing function (e.g. parsing sentence structures) and platform virtual machines, which are capable of supporting a wide variety of application VMs and possibly also new platform VMs.

One of the tasks for a designer is to specify the various forms of representation (types of syntax, or types of variability, for information-bearing structures) in which the information can be acquired, stored, manipulated, combined, derived, analysed, and used, as well specifying the ontologies that are capable of generating the required diversity of information contents, using those forms of representation.

Most researchers have experience of far too few forms. Many don't even realise this is a problem, either because they have never implemented a working system or because they have been educated into unthinkingly accepting only a narrow range of forms of representation as possible (e.g. vectors of numerical values and probability distributions over them, or perhaps purely logical forms, depending on their background).

In particular, I suspect we currently know very little about the forms of representation that allowed humans to investigate Euclidean geometry long before modern logic and algebra were developed, or the cartesian mapping between geometry and arithmetic. I have argued that it is likely that many animals have forms of representation whose instances are specially tailored to, but not isomorphic with, spatial structures and processes which they are used to represent.

(Pictures of impossible objects by Reutersvaard, Penrose and Escher give clues as to some of the reasons why isomorphism between what is represented and its representation is neither required nor desirable.)

f) Specifying conditions in which information about what is in the environment may or may not be available (easily, or with various kinds of difficulty) and the things that need to be done to obtain that information, using perception, inference, experimentation, communication, theory construction, etc. can be an important part of what would be requirements analysis for an engineer and specification of what needs to be explained, for a scientist.

Specifying the information-processing mechanisms that are capable of manipulating, storing, analysing, combining, transforming, and using the information, is often a highly creative design process for which empirical investigation of either visible behaviours or physical brain mechanisms can fail to yield vital clues.

- g) Without very varied design experience, researchers may consider only a subset of the purposes for which acquired or derived information can be used.

For an animal or robot, not only the kinds of factual information content about the environment mentioned above need to be represented. There also need to be formulations of information gaps (questions to be answered), goals, plans, preferences, values, hypotheses, experiments, and many kinds of control information, some of it transient in servo-control systems, others more abstract and enduring (e.g. about dangers to be avoided, preferences to be followed, etc.). (Compare theories about ventral and dorsal visual streams which fail to take account of these different functional requirements.)

Much of what needs to be controlled is not externally visible behaviour but internal information-processing behaviour, including controlling the perceptual interpretation of sensory inputs, controlling the competition between inconsistent goals or preferences, or sometimes selecting between planning options on the basis of unusual requirements for sensitivity to changing features of the environment during plan execution.

The design of mechanisms required for development and learning requires far more to be achieved than simply externally observable (and rewardable) changes in behaviour. Often there are deeper developments hidden from external view, including new forms of representation, new ontologies, new forms of storage and searching, new reasoning algorithms, new conflict resolution strategies.

Sometimes this requires changing the structure of an information-processing architecture, such as adding a whole new subsystem, or modifying connections between sub-systems. An example of such architectural revision seems to be the transition from pattern-based language use to grammar-based used, which causes children to start making errors because they don't yet have the ability to cope with exceptions to the grammatical rules -- which requires yet another architectural change. I suspect there are many more such architectural transitions that have gone unnoticed in developmental psychology (though Piaget collected evidence for some of them).

Very few researchers have personally encountered the diversity that has been explored in AI in the last half century because of the pernicious effects of factional warfare (or shortage of time) that leads senior researchers to teach only the types of mechanism and forms of representation that they happen to like or be familiar with. (The lack of substantial multi-disciplinary undergraduate degrees in AI/Cognitive science is partly to blame for this.)

- h) A problem for many researchers is understanding the possibilities for **development** of the forms of representation, algorithms, architectures, uses of information, and specific knowledge of the environment. Many AI researchers assume that the forms of representation available and the architecture are fixed from the start -- which clearly is not the case in humans, as Piaget noted a long time ago, though he lacked the conceptual tools required to express good explanatory theories.
- i) There are many unobvious designer tasks that have so far not received enough attention, including investigating **requirements** for organisms to acquire, manipulate, and use information about themselves, including information about their own location and relationships and processes in the spatial environment and also their own internal information processing -- i.e. treating the organism, including its internal information processing, as part of the environment to be perceived,

reasoned about and acted on. (Compare McCarthy on "Making robots conscious of their mental states", 1995, Minsky "Matter mind and models", (1968), Sloman 1978 Chapter 6).

Often it is assumed that requirements are clear and can be stated briefly, leaving only the problem of producing a design. When we are discussing systems whose functionality is the result of evolution over long time scales in which many different kinds and layers of functionality have been provided in a single design, the requirements that led to that design are generally far from obvious.

An example is the need to categorise different classes of use of information. Karen Adolf has referred to the "on-line intelligence" involved in infants and toddlers interacting with some complex and changing situation, such as walking across a narrow bridge, putting on clothes, pursuing escaping prey. That can be contrasted with "off-line" uses of information about an environment with which the individual is not currently interacting and may not soon interact. The off-line uses of information about the environment tend to go unnoticed by researchers who emphasise the importance of embodiment.

- j) I suggest that evident competences of both pre-verbal children and many non-human animals require the ability to make use of forms of representation that share some of the features previously thought to be unique to human language, including structural variability, varying complexity of forms, decomposability and recombining of parts, compositional semantics (particularly context-sensitive compositional semantics), and the ability to support inference.

Despite sharing those features with human communicative languages (and also logical, algebraic, and programming languages), these biologically useful forms of representation could be very unlike human language both in function, since their main use is not for communication between separate individuals, and in their form insofar as they are not all composed of linear sequences of discrete elements. (In the 1960s, for instance, some researchers on vision investigated "web grammars", for graph-like visual structures.)

As suggested above, there is reason to suppose that there is a deep relationship between the forms of manipulability of the information structures and the forms of manipulability of objects in space and time. Yet the representations cannot be isomorphic with what they represent: since many of the uses of information are very different from the uses of what is represented. (E.g. information about food and hiding places need not be edible, or provide shelter from rain.)

These considerations lead to new requirements for learning mechanisms that are capable of emulating the achievements of many intelligent mammals and birds and perhaps even some cephalopods. There are many details to be worked out and this remains a long term research project, but I hope some small steps in the new direction will be taken before long: there are probably already researchers unknown to me who are doing this.

NOTE

Some of the ideas presented above were inspired by discussions with Jackie Chappell and members of two EU-Funded cognitive robotics projects [The CoSy Project](#) and [The CogX Project](#).

The ideas are further developed in papers, presentations and discussion notes here:

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

References

(More to follow)

- A. Sloman, J. Chappell, 2005,
The Altricial-Precocial Spectrum for Robots,
Proceedings IJCAI'05, Edinburgh, pp. 1187--1192,
<http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>
- Aaron Sloman, 2008,
Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)?,
<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0702>
- Aaron Sloman, 2009,
Ontologies for baby animals and robots. From baby stuff to the world of adult science: Developmental AI from a Kantian viewpoint.
<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#brown>.
- Aaron Sloman, 2005
A (possibly?) new theory of vision.
<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk35>
- Jackie Chappell, Aaron Sloman, 2007,
Natural and artificial meta-configured altricial information-processing systems,
International Journal of Unconventional Computing, 3, 3, pp. 211--239,
<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>
- A. Sloman and J. Chappell, 2007
<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#wonac>
Presentations on Humean and Kantian conceptions of causation in animals and robots,
Presented at NSF/EU-funded Workshop on Natural and Artificial Cognition, Oxford, June 2007.
- Symposium on AI-Inspired Biology (AIIB) 2010
- To be added: refs to Gibson, Gibson and Pick, Karen Adolph, Kant, Gopnik, Tennenbaum, McCarthy, Piaget,

Maintained by [Aaron Sloman](#)
[School of Computer Science](#)
[The University of Birmingham](#)