

The Altricial-Precocial Spectrum for Robots *

Aaron Sloman

<http://www.cs.bham.ac.uk/~axs/>
School of Computer Science
University of Birmingham. UK

Jackie Chappell

<http://www.biosciences.bham.ac.uk/staff/staff.htm?ID=90>
School of Biosciences
University of Birmingham, UK

Abstract

Several high level methodological debates among AI researchers, linguists, psychologists and philosophers, appear to be endless, e.g. about the need for and nature of representations, about the role of symbolic processes, about embodiment, about situatedness, about whether symbol-grounding is needed, and about whether a robot needs any knowledge at birth or can start simply with a powerful learning mechanism. Consideration of the variety of capabilities and development patterns on the precocial-altricial spectrum in biological organisms will help us to see these debates in a new light.

1 Introduction

Attempts to design intelligent systems often encounter two objections: (a) for non-trivial systems the task is impossibly difficult, and (b) a system that has been designed is not intelligent since everything it does is a result of the designer's decisions, not its own. These and similar arguments lead some people to the conclusion that intelligent systems must start off without any knowledge, and acquire what knowledge and skills they need by interacting with the environment, possibly guided by a teacher, using only general learning capabilities.

This view assumes that something like human adult intelligence could be a product of neonates born with very little innate knowledge, but possessing a powerful general-purpose learning mechanism, e.g. reinforcement learning using positive and negative reinforcement based on drives, needs, and aversions. Then all knowledge would gradually be built up by continual shaping of internal and external responses to various combinations of internal and external stimuli. There have been several AI robotics projects and synthetic agent projects based on this line of thought (e.g. [Grand, 2004; McCauley, 2002; Weng, 2004], among others.)

This view ignores some powerful genetic influences in biological organisms. Most species show a mosaic of skills and abilities some of which develop mainly under genetic

influence, while others grow (usually gradually) as a result of learning and experience. Even learning mechanisms can have constraints tailored by natural selection to the probability of certain events co-occurring in their environment. For example, Domjan and Wilson [1972] showed that rats readily learned to associate the taste of water with sickness, and ringing bells with electric shock, but could not learn to associate noise with sickness or taste with shock.

Our ability to design artificial systems could benefit from a better understanding of the very varied spectrum of cases found in nature. At one extreme are biological species, often labelled 'precocial', whose young are born or hatched relatively well developed and able to cope with tasks that are far beyond what our current robots can do: newly hatched chickens look for food which they peck, and new-born deer can, within hours, find a nipple to feed from, and run with the herd. At another extreme (often labelled 'altricial'), animals are born or hatched helpless and underdeveloped, yet sometimes grow into adults with far more impressive cognitive skills than their precocial cousins, e.g. humans, hunting mammals and nest-building birds. How?

Precocial skills that appear in a relatively constant form in different environments can be fine-tuned by the individual experience of the animal. In contrast, altricial capabilities can vary widely within a species, as languages do between cultures, and as many kinds of competence do from one generation to another within a culture. Many children now have knowledge and skills beyond the wildest dreams of their parents at the same age. It is implausible that a combination of environmental differences and some simple general learning mechanism, or any known learning algorithms, could produce such variation. Evolution seems to have found something more powerful, which supports and enhances cultural evolution.

We propose that in addition to knowledge-free reinforcement learning mechanisms and the unlearned kinds of knowledge acquisition (like maturation, and peri-natally generated precocial skills), there are capabilities produced by powerful mechanisms characteristic of the altricial end of the spectrum, that explain how adult members of the same species can have hugely varying knowledge and skills. These mechanisms (a) acquire many discrete chunks of knowledge through play and exploratory behaviour which is not directly reinforced, and (b) combine such chunks in novel ways both in solving

*In Proceedings IJCAI'05, pp 1187–1192, Edinburgh, 2005
Many thanks to colleagues on the EC CoSy project.

problems and in further play and exploration. This process depends on a collection of very abstract ‘internal’ skills and knowledge with a strong innate component – though the collection can develop over time, especially in humans.

So evolution provides much information that animals do not have to learn: some used in overt behaviour in precocial skills and some in more abstract internal behaviour (e.g. skill-bootstrapping behaviour) towards the altricial end. We present a number of conjectures about how this happens and discuss implications for robot designers.

1.1 Precocial vs. altricial developmental mode

Conjecture 1: *There is no sharp altricial/precocial distinction between species, nor a continuum of cases, but a spectrum of capabilities, with many discontinuities. Each species exhibits a unique collection of skills from different parts of the spectrum. A similar variety of designs may be useful in robots.*

A number of factors determine the extent to which a behaviour or skill is modified by the animal’s own experience. The outcomes of these factors map very approximately on to the precocial and altricial developmental modes, though altricial species may have some precocial skills and *vice versa*. Analysis of these issues will prove valuable in deciding which sorts of robots will need ‘precocial’ capabilities, which ‘altricial’ capabilities, and what sorts of mechanisms are needed.

Factors promoting precocial skills and knowledge

The following will tend to favour evolution of precocial skills over altricial ones (if both are possible).

- **Opportunity:** If it would be very difficult for the animal to learn a skill given the constraints of its environment the skill has to be precocial. e.g. young blackcaps migrating for the first time on their own [Helbig, 1996], and the species-specific ‘password’ vocalisation in a brood parasite, the brown-headed cowbird [Hauber *et al.*, 2001], which is raised by other species.
- **Risk:** Learning depends on exposure to certain stimuli at the right time. If failing to perform a vital skill correctly would result in animals dying before reproduction, the skill is more likely to be precocial. e.g. predator recognition in Atlantic salmon, [Hawkins *et al.*, 2004], flight in cliff-nesting altricial birds.
- **Time constraints:** If learning would take too long to meet a need, a skill is more likely to be precocial. e.g. running in wildebeest calves, who must keep up with the herd soon after birth.

The young of species with many precocial skills start highly competent, requiring little or no parental care other than a degree of protection from predators.¹ They appear to have a great deal of information about structures in the world, how to perceive them, and how to behave in relation to them, even if they can be changed later in subtle ways as a result of adaptive processes, reinforcement learning or the like.

Since evolution can, and often does, provide huge amounts of information in the form of knowledge and skills. should not a large amount of precocial information be the norm, if all those millions of years of evolution are not to be wasted?

¹Note, however, that species with precocial young may have to invest more resources in their young pre-natally.

Factors promoting altricial skills and knowledge

In conditions opposite to those outlined above (abundant opportunity for learning, low risk outcome, no time constraints), altricial skills will tend to be favoured. In addition, the following factors are also important:

- **Unpredictability:** If the environment is very variable or unpredictable, there will be no consistent pattern for natural selection to act on. The fact that fleeing prey are relatively unpredictable, while grass tends to be highly predictable might explain why predatory mammals and birds tend to have altricial developmental patterns, while grazers have many precocial skills.
- **New niches:** This is related to unpredictability. If an animal encounters conditions which have never been encountered before in the evolutionary history of its species, individual learning will be the only possible response. Some niche variation can be accommodated by individual adaptation and calibration. But not when quite new concepts and new behavioural structures are required, e.g. if opponents often invent new weapons, or new defences against old weapons and strategies.

If altricial skills provide so much flexibility and diversity of responses, why don’t more animals show altricial skills? A partial answer is cost: neural tissue is expensive to build and maintain (in humans, the brain consumes about 20% of the basal metabolic rate). Learning can be time-consuming and risky. Young animals need to be protected from the harsh realities of life while acquiring altricial skills, so the intensity and duration of parental care can also be significant. The extended commitment of resources by parents to guarding, feeding, and in some cases helping their offspring to gain hunting skills, can both endanger the adults (make them more vulnerable to predation, deprive them of food given to the offspring), and reduce their breeding frequency.

1.2 The altricial-precocial spectrum for robots

Analysis of the altricial/precocial trade-offs already explored by evolution can provide a basis for a new understanding of the varied requirements and tradeoffs in robot design. Applications where tasks and environments are fairly static and machines need to be functional quickly, and where mistakes during learning could be disastrous (e.g. flight control) require precocial skills (with some adaptation and self-calibration), whereas others require altricial capabilities, e.g. where tasks vary widely and change in complex ways over time, and where machines need to learn to cope without being sent for re-programming (e.g. robots caring for the elderly). Architectures, mechanisms, forms of representation, and types of learning may differ sharply between the two extremes. And the end results of altricial learning by the same initial architecture may differ widely, as happens in humans.

1.3 Combining precocial and altricial skills

Precocial skills can provide sophisticated abilities at birth. Altricial capabilities add the potential to adapt to changing needs and opportunities. So it is not surprising that many species have both. However evolution seems to have discovered something deeper than just a mixture of some innate modules and some learning capabilities. Humans, and presumably some other animals, seem to have a *hybrid* architecture where the mechanisms for learning are

themselves a mixture of precocial and altricial mechanisms.

Conjecture 2: *Evolution ‘discovered’ and deployed the power of an architecture towards the more sophisticated, altricial, end of the spectrum, based not on specific knowledge about the environment, but on mechanisms using abstract knowledge about ways in which physically possible environments can differ. So powerful altricial learning arises from sophisticated genetically determined (i.e. precocial) learning mechanisms, which may be expanded during development.*

Instead of starting off able to recognise and react to specific stimuli, an altricial learner might have mechanisms that use knowledge about how to generate varied exploratory actions and how to seek recognisable patterns.² Such ‘playful’ actions reveal which among possible patterns are instantiated in the environment, and which ways of combining various kinds of pre-programmed or learnt action capabilities have effects that are ‘interesting’ not merely in terms of useful or harmful effects, but also according to criteria such as symmetry, elegance and power (e.g. producing big effects with little effort, or making something work that previously failed), thereby bootstrapping a wide variety of competences through interactions with the environment.³

2 What are the mechanisms?

This may sound like standard reinforcement learning, except for some innate biases to learn about certain classes of things only a subset of which are actually learnt under the influence of any particular environment [Cummins and Cummins, 2005]. But such mechanisms produce only gradual changes in behaviour, and it seems that at least some species, including humans, apes, and some corvids are — in addition — capable of kinds of discontinuous learning where hard problems are solved creatively by recombining previously acquired concepts, knowledge chunks and action chunks.

2.1 Gradual change vs learning chunks.

Conjecture 3: *Altricial learning is not always driven by biological needs and desires (e.g. for food, shelter, escape, mating, etc.) using rewards and punishments to drive gradual change through positive and negative feedback; instead the key feature is constant (often unmotivated) experimentation with external and internal actions during which re-usable chunks of information are learnt about what can be done and what can occur, and the preconditions and effects. Chunks are discovered both in passively perceived inputs (e.g. using self-organising classifiers to chunk sensory inputs) and also in behaviour patterns and their consequences.*

Re-use, in new combinations, of previously stored chunks, might explain the creative problem-solving of New Caledonian crows [Chappell and Kacelnik, 2004; Weir *et al.*, 2002]) and chimpanzees [Kohler, 1927], the ability of social animals to absorb a culture; and, in humans, the ability to learn and use a rich and highly structured language. A key feature of

such processes is learning new ontologies (discussed further below.)

2.2 Combinatorics: the origins of syntax

For these kinds of learning, mechanisms that gradually change weights or probabilities do not suffice: instead some chunks once learnt are explicitly referenced and available to be combined with other chunks, so as *immediately* to produce external and internal states and processes that are qualitatively (e.g. structurally) *novel* for the individual.⁴

Conjecture 4: *Some altricial species can not only store, label and categorise input and output ‘chunks’ that can be re-used later, but can also combine them to form larger chunks that are explored, and if found ‘interesting’ also stored, labelled, categorised, etc. so that they become available as new, larger units for future actions. A simple example would be picking something up and transferring something to one’s mouth, being combined to form a new unit to pick up an object and bring it to one’s mouth. Another example is combining moving right and moving left in a repeated sequence which can produce changes of visual input that provide useful information about 3-D structure [Philipona *et al.*, 2003].*

A simple demonstration of how this can enormously reduce search spaces, originally suggested by Oliver Selfridge over 20 years ago can be found in a tutorial program file provided with the Poplog system.⁵ In that example the ‘agent’ is given tasks which it attempts to solve using its primitive actions, previously learnt useful actions, and combinations of primitive and previously learnt actions formed by simple syntactic operations, including use of a ‘repeat construct’. If it is given a complex task initially it searches for a solution and gives up after a certain length of search. But as it gradually builds up useful re-usable chunks the space of actions in which it searches in a given number of steps gets more and more complex. If suitably trained by giving it the right problems to ‘play’ with it can suddenly become able to solve a class of quite difficult ‘counting’ problems.

This is a very simple example, but it illustrates a general point: that an altricial learning agent may need to be able to learn many things that are not intrinsically rewarding, but which are stored as a result of innately driven exploratory drives and then provide a basis for newer more complex explorations and learning – which could be specially well suited to the environment in which the simpler chunks (including perceptual and action chunks) have been learnt.

We have here a kind of syntactic competence which generates structures with compositional semantics.

Conjecture 5: *Not only external, but also internal actions – e.g. actions of goal formation, problem-solving, concept formation – may be composed, forming layers of complexity.*

If the methods of syntactic composition are themselves subject to the same process then the result may be production of ontological layers in many parts of the system including perceptual layering, action layering and various kinds of internal layering of control and description. Different kinds of competence may produce different kinds of layering.

²Manuela Viezzer’s PhD research on discovery of affordances illustrates this. See <http://www.cs.bham.ac.uk/~mxv>

³For more on interestingness see [Colton *et al.*, 2000]

⁴[Boden, 1990] distinguishes historical and personal creativity.

⁵<http://www.cs.bham.ac.uk/research/poplog/teach/finger>

2.3 Possible benefits of the hybrid design

There may be several different biological advantages gained by the hybrid architecture over the more obvious ‘precocial’ and ‘altricial’ architectures, which can be framed in terms of ways of overcoming resource limitations.

Advantage over precocial strategies

- The capabilities required in adults of some species may be too rich to be encoded in the genome if there is a ‘space’ limitation). (Genetic information capacity as a scarce resource)
- The evolutionary history of some species may not have provided any contexts in which certain currently useful capabilities could have been selected for, or may constrain the direction in which evolution can proceed (i.e. phylogenetic constraints) (Evolutionary history is a biological resource, and may have limitations like any other biological resource.)

Advantage over altricial strategies

- Although rapid, one-trial reinforcement learning is possible, reinforcement generally requires multiple exposures to a pairing of stimulus and reinforcer, and such slower forms of learning (involving gradual, adaptive, shaping, processes) may not be able to cope with the very varied environments and challenges facing adults of some species. (Time to learn as a limited resource)
- For some new problems an individual may have to produce a discontinuous change from previous behaviours, i.e. novel, creative behaviour, as opposed to interpolating or extrapolating in a space that is already well explored. That is because prior learning does not always take the individual ‘close’ to problems that can occur. (Individual learning opportunities as a limited resource)
- The benefits of precocial strategies may be outweighed by the risks of over-commitment in a particular direction. If there are constraints on types of learning that make certain sorts of developmentally fixed competence incompatible with kinds of learning e.g. because the prior commitment provides the wrong building-blocks for the required new competence and because it is difficult to re-wire or otherwise re-implement the required building blocks. In that case, the advantages of being precocial may be best abandoned in favour of the advantages of flexibility: provided that parents are available to care for helpless young. (Flexibility as a limited resource.)

All this points to the need for learning, decision-making, and acting capabilities that support discontinuous development either in the evolution of the species, or in learning by the individual, without adversely affecting fitness.

2.4 Effect of body and physical environment

Although innate altricial bootstrapping mechanisms encode a great deal of abstract know-how enabling specific kinds of development, the space of possible learnt concepts and action schemas permitted by such mechanisms (the ‘epigenetic landscape’) may be very large, leaving the environment scope to determine which subset of that space is actually learnt in both early and late development. Even within one environment, different kinds of bodies and different kinds of initial brain mechanisms will produce variation: (a) in the kinds of chunks that can be learnt, (b) in the number of different chunks that can be learnt (c) in the variety of forms of combination of previously learnt chunks.

Having two hands with five fingers each with several (more or less independently controllable) joints can produce huge combinatorial possibilities. Having grippers that can move independently of the eyes (i.e. hands vs beaks) can make a huge difference to perceivable patterns. Different designs for mouths (or beaks) and tongues or vocal mechanisms may also produce different combinatorial spaces to be explored.

A fixed learning mechanism may be able to acquire different sorts of reusable chunks and perhaps even different ways of combining chunks, when it is combined with different physical bodies, or different physical environments. Compare the different behavioural and conceptual results of exploratory play in children surrounded by rocks, sand, sticks, stones, mud and water, on the one hand and children surrounded by stackable blocks and cups, dolls, toy animals, meccano sets, lego sets, computers with mice and keyboards, on the other.⁶

Such learners start with both a control system that performs systematically varied internal and external operations driving these mechanisms, as well as mechanisms for learning, storing, labelling and re-using various kinds of chunks. (Compare macro-formation mechanisms in planners, e.g. HACKER [Sussman, 1975] and SOAR’s chunking.)

So the variety of sensory and motor chunks learnt will depend on (a) the physical design of the organism or robot, e.g. having independently movable limbs, digits, jaws, tongue, neck, head, eyeballs, ears, with motion changed in various ways, such as accelerating, decelerating, changing direction etc. (b) the richness of the organism’s environment, particularly during development (e.g. learning in rats improves with environmental enrichment [Iuvone *et al.*, 1996]). (c) the genetically determined collection of internal operations on data-structures, especially (d) the ‘syntactic’ mechanisms available for combining old chunks into new structures⁷ in perception (e.g. parsing), or in action (e.g. synthesising plans), and (e) whether those syntactic mechanisms can be applied not only to representations of external actions, but also to representations of internal actions such as operators used in problem formation, plan construction and debugging (as in Sussman’s HACKER).

So different altricial species will differ in many ways related to these differences. In particular some may have much richer combinatorial competence than others.

Some of the processes are clearly amenable to cultural influences: e.g. what sorts of toys and games are found in the environment during early development will make a difference to which ontologies can be acquired. The processes could lead to different kinds of understanding and representation of space, time, motion and causality in different species, and perhaps in different cultures within a species.

After mechanisms for chunking and composing evolved in forms that we share with some other mammals, further

⁶An example of such exploration can be seen in this video of a child aged 11 months exploring affordances in a world with yogurt: http://www.cs.bham.ac.uk/~axs/fig/josh23_0040.mpg

⁷Arp [2005] has recently made a similar point, but citing only the special case of visualisation. [Cummins and Cummins, 2005] propose only innate learning biases not these innate meta-level modules.

typical evolutionary processes of duplication plus differentiation could have produced variants that were particularly suited to human language, which, if used internally, would have enormously amplified all the effects discussed above. But our discussion implies that a powerful form of *internal* language developed earlier, to support processes of altricial bootstrapping that we appear to share with some non-verbal altricial species [Sloman, 1979].

3 Symbol ‘attachment’, not ‘grounding’


The existence of precocial species demonstrates that sophisticated visual and other apparatus need not be learned, implying that the semantic content of the information structures is somehow determined by pre-existing structures and how they are used. This eliminates, at least for such precocial species, one of the supposed explanations of where meaning comes from (“There is really only one viable route from sense to symbols: from the ground up” [Harnad, 1990]).

How then can innate abstract structures determine meaning, as we claim happens for precocial species? Our answer, presented in [Sloman, 1985; 1987], is inspired by ideas about models in meta-mathematics and proposals in philosophy of science about scientific terms that could not possibly get their meaning by abstraction from perceptual processes, (e.g. ‘electron’, ‘quark’, ‘neutrino’, ‘gene’, ‘economic inflation’, ‘grammar’).

A (consistent) formal system determines a non-empty class of possible (Tarskian) models, which may be either abstract mathematical objects or collections of objects, properties, relations, events, processes etc. in the world. For a given formal system not every such collection can be a model, though it may have many different models. As (non-redundant) axioms are added the set of possible models decreases. So the meaning becomes increasingly specific. A formal system can determine precise meanings for its theoretical terms, by ruling out the vast majority of possible things in the universe as referents.

There is typically residual ambiguity, but, as Carnap [1947] and other philosophers have shown, this residual indeterminacy of meaning may be eliminated, or substantially reduced, by the use of bridging rules or ‘meaning postulates’. As Strawson [1959] noted, further ambiguity in reference to individuals can in some cases be eliminated or reduced by causal relations between the event of use of symbols and the things referred to.

We can call this process ‘symbol attachment’, because it is a process in which a complex symbolic structure with rich structurally-determined meaning can be pinned down to a particular part of the world. This is very different from the theory of symbol grounding which requires all meaning to be derived bottom up by abstraction from sensory experience of instances, which is actually a very old philosophical theory (concept empiricism) put forward hundreds of years ago by empiricist philosophers and refuted by Kant. Our impression is that many of the AI researchers who now use the phrase ‘symbol grounding’ have ignored its original meaning (arising out of concept empiricism) and actually use it, unwittingly, to refer to something like what we have called



There are many intermediate cases between pure symbol grounding and extreme symbol-attachment. Arrows on the left indicate abstraction of concepts from instances or sub-classes. On the right arrows indicate a mixture of constraints that help to determine possible interpretations (models) of the whole system and/or make predictions that enable the system to be used or tested.

‘symbol attachment’. We suggest that in altricial species: (a) symbol attachment, in which structures produced by a combination of evolution, and exploratory ‘play’ and problem solving, acquire most of their meaning from their formal structure, with perceptual/action bridging rules to reduce their ambiguity, is far more important than what is normally called ‘symbol grounding’, and (b) learning of chunks and combinatory procedures tailored to the current environment but constrained by powerful innate bootstrapping mechanisms that evolved to cope with general features of the environment (e.g. concerning space, time, motion, causation, etc.) is far more powerful than most forms of learning now being studied.

4 Conclusion

The rapid, automatic, non-need-driven collection/creation of a store of labelled, reusable, perceptual and action chunks, along with ‘syntactic’ mechanisms for combinatorial extensions of those ‘basic’ chunks, provides a rich and extendable store of rapidly deployable cognitive resources, using mechanisms that will already sound familiar to many AI researchers who have worked on planning, problem solving, reasoning, 3-D vision, and language understanding, though we have said very little about how they might work.

It may be that if some of this happens while the brain is still growing the result can be ‘compiled’ into hardware structures that support further development in powerful ways, explaining why the process needs to start while the infant is still physically under-developed. (This is very vague.)

We are not saying that precocial species cannot learn and adapt but that the amount of variation of which they are capable is more restricted and the processes are much slower (e.g. adjusting weights in a neural net trained by reinforcement learning, as contrasted with constructing new re-usable components for combinatorially diverse perceptual and action mechanisms).

So we suggest that in addition to the experiments where we pre-design working systems with combinations of different sorts of competence in order to understand ways in which such competences might work and cooperate, we should also

begin the longer term exploration of architectures and mechanisms for a robot towards the altricial end of the spectrum described above, growing its architecture to fit opportunities, constraints and demands provided by the environment. Some further details regarding requirements for such architectures can be found in [Sloman *et al.*, 2005].

Readers with some knowledge of the history of philosophy will recognise all this as an elaboration of ideas put forward by I. Kant in his *Critique of Pure Reason* around 1780. He would have loved AI. Future tasks include relating this work to the ideas Piaget was trying to express, in his theories about a child's construction of reality, using inadequate conceptual tools (group theory, truth tables), and to the taxonomy of types of cognitive evolution in [Heyes, 2003].

Application domains where tasks and environments are fairly static and machines need to be functional quickly, require precocial skills (possibly including some adaptation and self-calibration), whereas others require altricial capabilities, e.g. where tasks and environments vary widely and change in complex ways over time, and where machines need to learn how to cope without being sent for re-programming. Architectures, mechanisms, forms of representation, and types of learning may differ sharply between the two extremes. And the end results of altricial learning by the same initial architecture may differ widely.

If all this is correct, it seems that after evolution discovered how to make physical bodies that grow themselves, it discovered how to make virtual machines that grow themselves. Researchers attempting to design human-like, chimp-like or crow-like intelligent robots will need to understand how. Whether computers as we know them can provide the infrastructure for such systems is a separate question.

References

- [Arp, 2005] Robert Arp. Scenario visualization: One explanation of creative problem solving. *Journal of Consciousness Studies*, 12(3):31–60, 2005.
- [Boden, 1990] M. A. Boden. *The Creative Mind: Myths and Mechanisms*. Weidenfeld & Nicolson, London, 1990.
- [Carnap, 1947] R. Carnap. *Meaning and necessity: a study in semantics and modal logic*. Chicago University Press, Chicago, 1947.
- [Chappell and Kacelnik, 2004] J Chappell and A Kacelnik. New Caledonian crows manufacture tools with a suitable diameter for a novel task. *Animal Cognition*, 7:121–127, 2004.
- [Colton *et al.*, 2000] Simon Colton, Alan Bundy, and Toby Walsh. On the notion of interestingness in automated mathematical discovery. *Int. Journ. of Human-Computer Studies*, 53(3):351–375, 2000.
- [Cummins and Cummins, 2005] D Cummins and R Cummins. Innate modules vs innate learning biases. *Cognitive Processing: International Quarterly Journal of Cognitive Science*, 3(3-4):19–30, 2005.
- [Domjan and Wilson, 1972] M Domjan and N E Wilson. Specificity of cue to consequence in aversion learning in the rat. *Psychonomic Science*, 26:143–145, 1972.
- [Grand, 2004] Steve Grand. *Growing Up with Lucy: How to Build an Android in Twenty Easy Steps*. Weidenfeld & Nicolson, London, 2004.
- [Harnad, 1990] S. Harnad. The Symbol Grounding Problem. *Physica D*, 42:335–346, 1990.
- [Hauber *et al.*, 2001] M E Hauber, S A Russo, and P W Sherman. A password for species recognition in a brood-parasitic bird. *Proceedings of the Royal Society of London B*, 268:1041–1048, 2001.
- [Hawkins *et al.*, 2004] L A Hawkins, A E Magurran, and J D Armstrong. Innate predator recognition in newly-hatched atlantic salmon. *Behaviour*, 141:1249–1262, 2004.
- [Helbig, 1996] A J Helbig. Genetic basis, mode of inheritance and evolutionary changes of migratory directions in palearctic warblers (aves: Sylviidae). *Journal of Experimental Biology*, 199:49–55, 1996.
- [Heyes, 2003] C.M. Heyes. Four routes of cognitive evolution. *Psychological Review*, 110(713-727), 2003.
- [Iuvone *et al.*, 1996] L Iuvone, M C Geloso, and E Dell'Anna. Changes in open field behavior, spatial memory, and hippocampal parvalbumin immunoreactivity following enrichment in rats exposed to neonatal anoxia. *Experimental Neurology*, 139:25–33, 1996.
- [Kohler, 1927] W. Kohler. *The Mentality Of Apes*. Routledge & Kegan Paul, London, 1927. 2nd edition.
- [McCauley, 2002] L. McCauley. *Neural Schemas: Toward a Comprehensive Mechanism of Mind*. PhD thesis, University of Memphis, 2002. <http://csrg.cs.memphis.edu/mccauley/McCauley-Dis.pdf>.
- [Philipona *et al.*, 2003] D. Philipona, J.K.O'Regan, and J.-P. Nadal. Is there something out there? Inferring space from sensorimotor dependencies. *Neural Computation*, 15(9), 2003. <http://nivea.psych.univ-paris5.fr/Philipona/space.pdf>.
- [Sloman *et al.*, 2005] A. Sloman, R.L. Chrisley, and M. Scheutz. The architectural basis of affective states and processes. In M. Arbib and J-M. Fellous, editors, *Who Needs Emotions?: The Brain Meets the Machine*. Oxford University Press, Oxford, New York, 2005. Online at <http://www.cs.bham.ac.uk/research/cogaff/0-INDEX03.html#200305>.
- [Sloman, 1979] A. Sloman. The primacy of non-communicative language. In M. MacCafferty and K. Gray, editors, *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979*, pages 1–15, London, 1979. Aslib. Online at <http://www.cs.bham.ac.uk/research/cogaff/>.
- [Sloman, 1985] A. Sloman. What enables a machine to understand? In *Proc 9th IJCAI*, pages 995–1001, Los Angeles, 1985.
- [Sloman, 1987] A. Sloman. Reference without causal links. In J.B.H. du Boulay, D.Hogg, and L.Steels, editors, *Advances in Artificial Intelligence - II*, pages 369–381. North Holland, Dordrecht, 1987.
- [Strawson, 1959] P. F. Strawson. *Individuals: An essay in descriptive metaphysics*. Methuen, London, 1959.
- [Sussman, 1975] G.J. Sussman. *A computational model of skill acquisition*. American Elsevier, 1975.
- [Weir *et al.*, 2002] A A S Weir, J Chappell, and A Kacelnik. Shaping of hooks in New Caledonian crows. *Science*, 297(9 August 2002):981, 2002.
- [Weng, 2004] J. Weng. SAIL and Dav Developmental Robot Projects, 2004. <http://www.cse.msu.edu/~weng/research/LM.html>.