How a Philosopher Became an Information-Scientist
Answers to Luciano Floridi's Five Questions

(For a book to be published by Automatic Press/VIP:

`http://www.vince-inc.com/automatic.html`)

Last updated: June 30, 2008

Name : Aaron Sloman
Web: `http://www.cs.bham.ac.uk/~axs/`
Position :
Honorary Professor of Artificial Intelligence and Cognitive Science
School of Computer Science, The University of Birmingham, UK

Contents

1 How did it start?

Question 1: Why were you initially drawn to computational and/or informational issues?

I shall first give a high level, shallow overview of the answer, and then a more detailed answer, providing some of the substance.

1.1 High level overview

The move came originally from my concern about spatial reasoning, which I felt was at the core of important kinds of mathematical reasoning and many other kinds of human capability. That was the subject of my DPhil thesis (Sloman, 1962), which attempted to defend some of Immanuel Kant's ideas about synthetic necessary truths (Kant, 1781).

My first encounter with computational models, in 1969, resulted from the arrival of Max Clowes, an AI vision researcher, in the Experimental Psychology laboratory at Sussex University.[1] Our discussions (which also included Keith Oatley) convinced me that in order to understand key features of human spatial reasoning, and the spatial reasoning used by some animals, we need to adopt the design stance[2] and understand the various ways in which perceptual systems and other systems that process information can work. That understanding is accelerated by attempting to design and build machines that do such things, which provides a much deeper understanding of the design options and the tradeoffs between them than philosophers can get from arm-chair discussion and speculation.

It soon became clear to me that vision could not be understood in isolation: visual mechanisms work in combination with many others within a complete functioning *architecture*. For example, my interest in vision was originally triggered by its role in mathematical reasoning. Moreover, it was clear that vision not only provided factual information about the environment but also had aesthetic functions, attention control functions, action-control functions, sexual functions, and a variety of different sorts of communication function, including reading facial expressions and gestures, reading text, reading diagrams, and reading music. So early on I began to think about integrated information-processing architectures combining many different varieties of components, and that eventually led me to the design-based analysis of many other aspects of human minds and animal minds, constantly driven, by the question: what sort of machine could do *that*?

[1] In principle I could have learnt a great deal from my colleague Margaret Boden, who, by then, had been reading about and writing about AI. But for some reason we did not talk about that topic until much later, when I discovered that she already knew about most of the major developments in computational cognitive science, as revealed in her second book (Boden, 1978), which became one of the leading introductory AI textbooks, especially for people in other disciplines, because of her outstanding ability to explain (and criticise) the important developments in an implementation-independent way, demonstrated also in (Boden, 2006).

[2] As explained here
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/design-based-approach.html`

Some of the ideas, including the importance of information-processing *architectures* were reported in my 1978 book (written at the behest of Margaret Boden), in which I tried to show (in chapter 2) how the enterprise, like theoretical linguistics, required a broadening of popular conceptions of the nature of science, as well as showing in other sections how it changed our ideas about the nature of mind and relations between mind and body.

The more I learnt, the more I realised how hard the problems were, and how impressive the achievements of evolution were. There are still many unsolved problems requiring interdisciplinary cooperation, and when I can find philosophers, psychologists, neuroscientists, biologists, roboticists, linguists, computer scientists, etc. interested in addressing the issues, I try to learn from and work with them, though they are rare, partly because our educational and research funding systems produce and reward specialists with narrow vision, and also because the educational opportunities promised three decades ago by the development of computing, namely making it possible for many more people to learn how to think about complex working systems were not taken up. The failure, which is one of the great missed opportunities of recent decades, is discussed further in Section 1.2.5.

1.2 A longer version of the story

1.2.1 DPhil Research

My Oxford DPhil thesis in 1962 (now online) was about why Kant was right about mathematical discovery and Hume (and most contemporary analytic philosophers) were wrong. I tried to show why some kinds of mathematical discovery were both substantive (synthetic, in Kant's terminology) and non-empirical (i.e. a priori), for example geometric theorems, originally proved using the human ability to visualise geometric constructions and their consequences.

I kept on thinking about the problems after finishing the DPhil. For example, in Sloman (1965b) I tried to explain some of the distinctions required for understanding Kant's view (e.g. the necessary/contingent, apriori/empirical and analytic/synthetic distinctions) and in Sloman (1965a) extended some of Frege's ideas about functions in an attempt to analyse some sources of necessity. I returned to examining the notion of seeing that something had to be true in Sloman (1968/9.), but I was not really happy with what I was able to say about our ability to "see" necessary structural relationships. Later I realised that I needed new conceptual tools, which AI seemed to be developing.

1.2.2 Meeting Max Clowes

Around 1969 I met Max Clowes, a well known highly charismatic AI vision researcher (Sloman, 1984), who had joined the Experimental Psychology department. Max and I had many discussions and I attended the lectures on programming and AI that he presented to psychology students. The work he and others were doing on interpreting images, used only logical formulae (or equivalent symbolic structures) to express the interpretations. This did not address requirements to keep the interpretation in registration with image features, so that the information produced could be described as “visual” and could be operated on by using its spatial structure. I told him I thought that could be done by adding symbolic information to the image structures. This would also facilitate geometric modes of reasoning about the scenes, which I was sure were needed for many purposes besides visual perception, including geometric reasoning in mathematics.

In 1971, when Max was one of the organisers of the 2nd International Joint Conference on AI at Imperial College London, he persuaded me to submit a paper on my ideas about geometrical reasoning, so I wrote Sloman (1971), attacking the logicist methodology propounded by McCarthy and Hayes (1969), which Max had drawn to my attention. I made a distinction between Fregean and analogical representations, focusing mainly on their role in reasoning. This was in part an attempt to clarify the frequently made distinction between symbolic and pictorial or imagistic representations, often based on the mistaken view that pictures and diagrams are isomorphic with what they represent, clearly refuted by 2-D pictures of 3-D objects. Another closely related common mistake was to think the distinction was about the difference between discrete (or digital) and continuous information structures. Likewise, most people were unaware of the important features of symbolic representations that had been analysed by Frege, namely that they produced more complex structures from simpler ones by (recursively) applying functions of various sorts to arguments of various sorts, and making use of compositional semantics. Sloman (1965a) was an attempt to generalise some of Frege’s ideas regarding compositional semantics, and I have only recently realised that a similar generalisation can be applied to analogical representations.[3]

My paper argued that the relationship between an analogical representation and what it represents is far more subtle than most people had realised, and, as I had learnt from the work of Clowes and others, can involve highly context-sensitive modes of interpretation that vary from one part of an image to another, for instance in a 2-D picture showing several surfaces with different 3-D orientations. The paper also pointed out that reasoning by modifying a spatial structure (changing either geometric or topological features) had much in common with reasoning by modifying logical or algebraic formulae, though the latter is much easier to implement on a computer.

I realised then that if we could build a working model of a human-like visual system we might be able to use it as the basis for a working model of human geometric reasoning in both mathematics and in every day life (e.g. causal reasoning), making a major contribution both to philosophy of mathematics and possibly also mathematics education, though I did not have the programming experience required for that.

1.2.3 A formative year in Edinburgh

Fortunately, after reading my 1971 paper, Bernard Meltzer, head of the Computational Logic department in Edinburgh University, and founding editor of the journal *Artificial Intelligence*, in which he published the paper, obtained a grant from the UK Science research council to bring me to Edinburgh for a year in 1972-3. Because there were so many outstanding AI researchers

[3] Explained in `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang`

either based in Edinburgh or visiting during the year, that was an amazingly educative year for me. Among other things, it was the year in which the Edinburgh AI team got their robot Freddy to assemble a simple toy car and a toy boat, using 3-D vision and a robot hand, a very impressive achievement that has largely been ignored by recent robotic researchers.[4] It was while in Edinburgh that I first learnt about the importance of virtual machines, partly as a result of reading Wegner (1968), and partly as a result of experiencing the benefits of operating systems with virtual memory mechanisms. The Edinburgh POP-2 system on which I wrote my first programs supported up to 8 simultaneous users and allocated each user a virtual address space, which was frequently re-mapped onto the underlying hardware by the garbage collector, in order to optimise use of the very scarce physical memory in the computer (128 Kbytes!). This made it obvious that facts about supervenience in computer virtual machines contradicted widely held philosophical assumptions about supervenience.

1.2.4 Working on vision at Sussex and Birmingham

Soon after returning to Sussex I managed to get my own research council grant for three years, during which I had David Owen and later Geoffrey Hinton as research fellows on a project to test some ideas about the architecture of a human-like vision system. Some of the results of that work are reported in the chapter on vision in (Sloman, 1978a), describing the Popeye vision system which demonstrated how a mixture of *concurrent* top-down, bottom-up and middle-out processing at different ontological levels could enable rather complex and noisy images to be interpreted, using both the data in the images and prior knowledge about possible contents of the world. This combination of concurrent collaborative computation and prior knowledge explained both the speed of recognition in easy cases and graceful degradation as images grew more complex, features that are often thought to be impossible for systems implemented on computers.

However it was then clear that we were nowhere near a full vision system, and as I continued working on requirements for designing a working vision system over many years,[5] I constantly felt that the problems of explaining human vision were far more complex and varied than most vision researchers in AI, psychology or neuroscience (with the possible exception of Trehub, (Trehub, 1991)) had realised, partly because of the multiplicity of forms of representation involved, and the multiplicity of ontologies involved in perceiving structures and processes in the 3-D environment, and partly because of the multiplicity of types of process that could occur in a visual system, e.g. seeing spatial features, spatial structures, spatial structural relations, causal relations, affordances, processes, gestures, facial expressions, social interactions and, as a result of various kinds of learning processes, being able to read written languages, sightread music, read mathematical and programming notations, and reason about mathematical problems. (That is nowhere near a complete list. In particular it omits all the aesthetic functions of vision.)

As pointed out in Sloman (1982), some of these visual abilities involved acquiring and using

[4]See `http://www.aiai.ed.ac.uk/project/freddy/`

[5]See Sloman, 1982, 1989, 1994, 1993, 1995, 1996, 1998, 2002b, 2001, 2006.
A very recent example is a discussion paper on predicting affordance changes using spatial reasoning:
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0702`

transiently available information about continuously varying processes, information that was lost immediately after use, whereas other visual abilities involved acquiring and storing (at least for a while) *re-usable* information about objects, surfaces, relationships, and constraints that could play a role in reasoning (including planning actions and predicting what was going to happen in the environment) and in communicating information about the environment to others. For instance, in continuous visual control of actions, such as running a finger along a curved line, much transient information is used in making continuous adjustments to the motion, information that is lost after it has been used, unless the information-processing architecture allows such processes to be explicitly monitored and information about what they do recorded so that it is available for future use, for instance in future deliberative reasoning. This requires the information to be discretised into categories useful for forming generalisations.[6] This goes beyond the functionality required for performing the actions, though it can support learning that improves how actions are performed. Incomplete understanding of the different ways in which information can be acquired, represented and used in different parts of a multi-functional architecture has led to much confusion among researchers about dual (dorsal and ventral) visual streams and so-called "mirror neurons". If more people tried to design working models inappropriate theories might not survive so long.

1.2.5 Growing COGS at Sussex

I feel it is necessary to include a section on our work on teaching, because a great deal of what I learnt in the 1970s and 1980s about computing and AI and their relevance to other disciplines, arose from the intense collaboration with colleagues at Sussex in developing new teaching ideas and tools to support the learning. These tools and the related philosophy were, for a while, used in a number of other university departments, and in one UK school (Marlborough College). They also formed the basis of the new AI half degree introduced in Birmingham after I moved there.

After my return to Sussex, a team of colleagues, including Margaret Boden, Max Clowes, then later Steve Hardy, Gerald Gazdar and others, started an undergraduate programme in the Arts and Social Sciences area of the University, in which we taught AI alongside philosophy, linguistics, psychology, and at first anthropology, though that was later dropped as none of the local anthropologists was doing research in the area of overlap. The teaching was highly collaborative and was closely related to our research. A few years later we introduced an undergraduate AI major and also a 'conversion' MSc degree in 'Knowledge Based Systems'. Eventually, this combined activity – undergraduate teaching, MSc teaching, PhD supervision, and research – grew into COGS, The School of Cognitive and Computing Sciences, with Margaret Boden as first dean. Those were exciting times, with much wide-ranging discussion which helped to extend my ideas though I can't recall details of what I learnt from whom: leading figures with a shared belief that AI could be the glue that held several disciplines together included Max Clowes, Margaret Boden, Steve Hardy, Gerald Gazdar, Chris Mellish, Jonathan Cunningham, and later on David Hogg and David Young, with Steve Isard helping from his base in the Experimental Psychology group, where he taught AI and shared in the development of our teaching materials.

[6] Some of the requirements are discussed in more detail in this discussion of "fully deliberative" architectures: `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604`

Christopher Longuet-Higgins was also there but had only research commitments. We had a stream of distinguished visiting researchers. Alan Mackworth, whose D.Phil was supervised by Max Clowes also contributed in the early years.

Initially we taught AI by using Algol-68 in batch mode on the university central mainframe computer. Frank O'Gorman who then worked for Max Clowes taught me an enormous amount about Algol-68, a sophisticated language that stretched many of my ideas about computation, but was unsuitable for beginner AI students. So we tried using a batch mode version of Pop-2. But it soon became clear that for AI students development and testing had to be interactive: one could not simply develop (or read up) a theory, work out its mathematical formulation, then type in the code with some test input, as some researchers in other disciplines were accustomed to doing. Testing and debugging game playing, or language understanding, or incremental learning programs, needed to be highly interactive. Fortunately, from around 1975 the university gave us funds to buy our own PDP11 computer. We decided that none of the existing readily available AI systems for that machine would meet our needs, so Steve Hardy developed an implementation of the Edinburgh AI programming language POP-2, which he called Pop-11. Although he was the programming expert, several of us were involved in designing features of the language, including Max Clowes and Steve Isard. As a result, Pop-11 extended Pop-2 in several ways, mainly to support AI teaching for students whose background was not necessarily scientific or mathematical. Some of those extensions also made the language more useful for our research, especially the inclusion of a pattern matcher as part of the language.

Contributing to the design and implementation of Pop-11 and its suite of teaching libraries for AI, helped me acquire a much deeper understanding of the nature of computation than I could have had merely as a user of a programming language. This work continued during the 1980s as Pop-11 grew into Poplog, a multi-language development environment with a number of novel features, mostly designed and implemented by John Gibson, who became the chief architect. When it turned out that people in industry wanted to use it we were able to get funds from sales and research grants to help with the development, and eventually Poplog became a successful product, which helped many UK researchers in academe and industry learn about and use AI programming techniques.

Partly because of my own learning experiences, and partly on the basis of our teaching experiences from 1976 onwards, I became convinced that the teaching of programming, and especially teaching children to design, implement, test, debug, document, analyse and compare AI programs, would have a profound effect on the development of education and could lead to important changes in ways humans understood and thought about structures and processes, including mental structures and processes. Some of those educational ideas were presented in (Sloman, 1978a). and since then I have tried to document them in online web sites.[7] Unfortunately, this potential was never realised because politicians, educators, teachers and well-meaning parents made the deep mistake of thinking that the most important thing to teach children about computing was how to use the tools they were likely to have to use in their jobs, e.g. word processors, databases, and later web-browsers, etc., instead of how to design, implement, debug, analyse and

[7] E.g. in these two `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/compedu.html`
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/dscedu.html`

compare working systems. I often wonder how many politicians who waste tax-payer's money on unworkable schemes (like the UK's national identity card scheme, and the national health service IT project[8]) might have been more cautious if they had had real experience of trying to make moderately complex systems work.

This enormous missed educational opportunity has mostly gone unremarked (though its consequences are often lamented, without the causes being recognized). Recently, an influential computer scientist (Wing, 2006) has noticed the problem, and has been eloquently recommending a change of direction, but it may now be too late to undo the damage.

The experience of managing the Poplog development team and interacting with the commercial distributors (initially SDL in 1983, and later a start-up company ISL, formed in 1989 by a subset of the main company) also taught me a great deal about AI and its applications. That process continued until I moved to the School of Computer Science at the University of Birmingham in 1991, though I never concealed the fact that I was primarily a philosopher, doing philosophy in a new way, not a computer scientist.

1.2.6 Working on robotics at Birmingham

My ideas about architectures, and especially the relationships between architectures and affective states and processes continued to develop in Birmingham, at first in collaboration with Glyn Humphreys in psychology. I was very lucky to have some excellent PhD students who stretched my mind. For example, Luc Beaudoin, who was investigating architectural requirements for emotions, made me realise that the ideas about concurrency in chapter 6 of Sloman (1978a) had to be extended by decomposing some of the higher level processes into concurrent management (e.g. deliberation and problem solving) and meta-management (i.e. monitoring and modifying the management processes, and others), and who also made me realise that a typical decision-tree is not suitable for such meta-management processes, because too many of the sub-problems are interdependent and can interact in ways that an algorithm designer could not anticipate. I have not yet seen any proposed architecture that deals with all the requirements. (Minsky's (2006) comes closer than others I know of.)

For a while my research on vision took a back seat, though I made a number of unsuccessful attempts to get funding for a project to work on a real or simulated robot to explore the relationships between 3-D vision and action, among other things.

My work on those topics progressed slowly, but eventually accelerated as a result of working in a collaborative EU-funded robotic project (`http://www.cs.bham.ac.uk/research/projects/cosy/` which started in 2004). Detailed work on requirements for the architecture and representations used by a human-like robot performing everyday domestic tasks, led both to the realisation that perceptual contents are primarily about *processes* of many sorts (i.e. perceived structures in the environment are really perceived processes with little or no change) and that understanding vision and understanding causation in the physical environment were

[8]Whose folly is explained here `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/isoft/`

closely related.[9] Further progress followed the arrival in Birmingham of Jackie Chappell, who studies animal cognition, especially bird cognition, in the School of Biosciences. One example of our collaboration is a linked pair of workshop presentations on understanding Humean and Kantian causation: `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac`. We have also begun to develop new theories about tradeoffs between evolution of cognition and development of cognition (e.g. in Sloman & Chappell, 2005, and Chappell & Sloman, 2007).

Most of the problems are still unsolved: we do not have working visual systems remotely like those of birds or mammals in their capabilities, and a deep model of human spatial understanding and reasoning still seems a long way off (though simple fragments have been modelled in connection with mathematical reasoning in Jamnik, Bundy, and Green (1999) and Winterstein (2005)). We also still lack a philosophical theory of causation, rich enough to be the basis of an implementation of a machine able to engage in Kantian causal reasoning, though there is progress in essentially Humean causal reasoning, based on Bayesian reasoning. We also lack a theory of human learning, though there are many fragments available from many disciplines, most of them still ignored in AI, especially lessons from biology. There has, of course, been a very influential fashion for exploring and attempting replicate biological mechanisms. The mistake is to ignore what the mechanisms are required for. This leads to over-simplified tasks and benchmarks for testing the models.

1.2.7 Evolution, development and GLs

Many people think that to explain human intelligence we must show how an animal or robot with no prior knowledge about the environment, or itself, could use a very general learning capability (such as reinforcement learning), possibly aided by teachers, to acquire all the knowledge it will ever need. This idea has been implausible as regards humans for many decades, and, as McCarthy remarked, “Evolution solved a different problem than that of starting a baby with no a priori assumptions” (McCarthy, 1996).

Many years ago when I learnt that many animals, members of so-called ‘precocial’ species, show extraordinary competences soon after birth, or hatching, such as young grazing mammals that get up and walk to the mother’s nipple, deer that run with the herd, chicks that feed themselves and follow the nearest hen they see after hatching, and many more, I took that as evidence that evolution is capable of producing extremely sophisticated and knowledgeable neonates (apparently with much richer innate knowledge than either Descartes or Kant ever assumed possible).

However, that raised the question: Why do some species, whose adults seem to be particularly intelligent, including nest-building birds, hunting mammals, and primates, especially humans, apparently start off at a much less knowledgeable level? This led to a hunch that the starting point for such ‘altricial’ animals was actually more sophisticated than it appeared to be, and included mechanisms tailored to our environment in a generic manner, which provided potential for more advanced development than the other species that appeared more intelligent at birth.

[9] These ideas are developed and related to the notion of perceiving “proto-affordances” in Sloman (2008a).

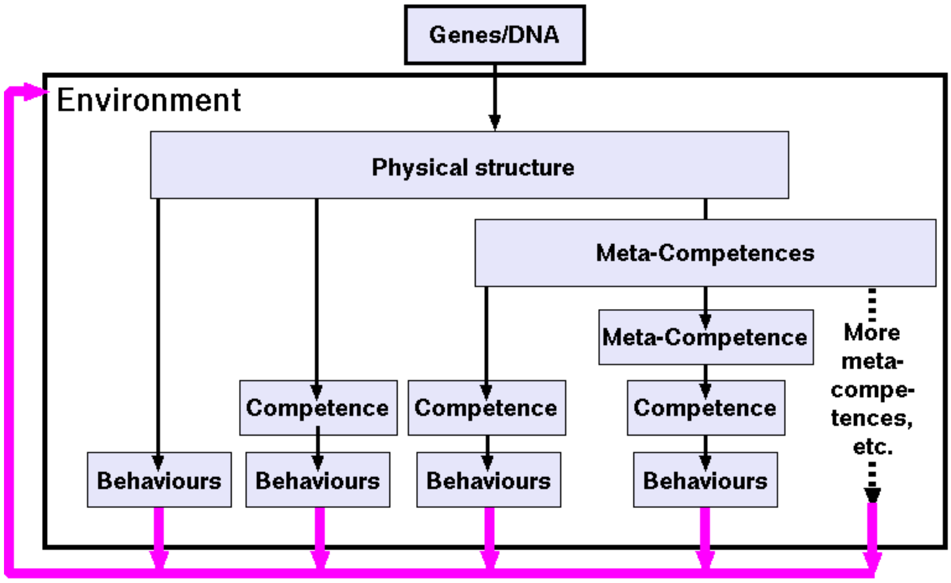

(Diagram from Chappell and Sloman (2007), produced with help from Chris Miall).

Figure 1: *This illustrates multiple routes from the genome to behaviours. The routes to the left of the diagram are mostly genetically determined, but the behaviours they produce and the environment's responses combine with genetically determined and acquired meta-competences to produce new competences (including ontology extensions) which generate routes to behaviour further to the right of the diagram. This cognitive epigenetic process includes growth of an information-processing architecture which is the joint product of the genome and, at different levels of abstraction, the environment.*

Ever since then I have been trying to characterise that difference in information-processing terms. The work gained a new impetus when Jackie Chappell joined in, as explained above. We have several papers developing ideas about the precocial-altricial (or in our preferred terminology, preconfigured-metaconfigured) spectrum for types of competence, including partial specifications for bootstrapping mechanisms capable of rapidly exploring and learning about the environment on the basis of innate meta-knowledge about what sorts of thing can be learnt and how they can be learnt, as opposed to specific knowledge such as precocial species display. Figure 1 summarises some of the ideas we have been developing about routes from the genome to behaviours.

Recently we have argued (reviving some themes from Sloman, 1978b and Sloman, 1979) that the kinds of learning and kinds of competence shown by many animals that lack the ability to communicate in a language with a rich syntax would be impossible if they did not use a type of 'Generalised Language' (GL) *internally* for expressing information that can be complex, highly structured, and very variable. Similar comments apply to pre-linguistic children. More specifically, such languages must support expression of information contents that vary in structure and complexity (unlike the fixed size vectors typical of inputs and outputs of many neural net models) and which support context sensitive compositional semantics. This internal generalised language need not be restricted to combinations of discrete atomic symbols, and may in some cases include geometrical structures used to express meaning and inferences (as emphasised in the 1971 paper).

We conjecture that sophisticated GLs evolved before human language, and are used by some non-human animals, and that young children use complex GLs for internal information-processing long before they can use an external language to converse in. This has implications for both the evolution and the development of human language, on which we have a very brief published paper (Sloman & Chappell, 2007) and an online presentation.[10] Some conjectured mechanisms that may be involved in linguistic and other forms of cognitive development, including development of several layers of ontology while learning about the environment, are sketched at a high level of abstraction in (Chappell & Sloman, 2007). There is still much to be done to develop the ideas, and we are hoping that experiments in robotics as well as further research on animal cognition and human learning will shed light on the mechanisms and processes involved. The kinds of innate competence implied by such theories must have many implications for biological as well as robotic research. A good source of information about some of the facts that such mechanisms need to be able to explain is Gibson and Pick (2000). Adequate working models could shed new light on Kant's theories about synthetic *a priori* knowledge, in part by explaining how mathematical competences develop.

1.2.8 The central importance of architectures

At a fairly early stage it became clear to me that trying to understand vision on its own was impossible. Likewise other kinds of mental function. In humans and many other animals mental capabilities form a complex integrated system with many components coexisting and interacting concurrently with one another and, in some cases with the environment. This led me to argue that we need to think not just in terms of algorithms and representations, as was then commonplace in AI, but in terms of *architectures*. I wrote a paper on this during my stay in Edinburgh in 1972-3 and it became a departmental discussion paper (Memo 59 Department of Computational Logic, slightly revised and published in 1973 in the *AISB Newsletter*. A still later version became chapter 6 of Sloman (1978a).

It also became clear that the notion of architecture was also very important for philosophy. That is because philosophers mostly attempt to analyse concepts related to human mental competences, such as belief, desire, emotion, imagination, learning, memory, creativity, and consciousness, as if these concepts all had clear meanings and "correct" definitions.

However, there are three important facts that impinge on this:

- different animals have different information-processing architectures, and humans have different architectures at different stages of development or in some cases because of brain damage or disease
- different architectures make very different states and processes possible

[10] `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang`

- concepts of mind refer to states and processes in complex information-processing systems.

It follows that not all mental concepts are applicable to all architectures. This means that instead of philosophers attempting to produce unique and correct analyses for concepts of the sort illustrated above, they should map out ways in which those concepts can be related to different architectures, including allowing that some architectures make certain concepts totally inapplicable.

For example, if a system has a collection of programs that can be triggered into action by different sensor inputs, and if when each program runs nothing else is happening, then that system is incapable of monitoring what it is doing or deciding to modulate what it is doing. So it cannot be described as self aware or self controlled. [11] For philosophers to explore the space of possible architectures and the different families of concepts that are applicable to different sorts of architecture they need to learn to think like designers.

1.2.9 Thinking like a designer about emotions and other forms of affect

Some of the ideas about architectures underpinning various kinds of mental states and processes were further developed after I moved to Birmingham in 1991, and worked with a series of PhD students, initially in collaboration with the psychologist Glyn Humphreys, on trying to clarify notions like, desire, motive, intention, preference, emotion, mood, attention, by showing how they could relate to processes that occur and dispositions that arise in an architecture that we conjectured was typical of a normal adult human being. In particular, Luc Beaudoin, Ian Wright and I tried to show how many of the phenomena of long term emotions, such as grief, jealousy, infatuation, obsessive ambition, and others, which can endure even when they are not being manifested or experienced, though they may easily be triggered into expressing themselves, could be accounted for within the type of architecture we were developing, as reported at length in Wright, Sloman, and Beaudoin (1996).

We proposed an explanation of some of the variety of types of emotion, including enduring but sometimes dormant emotions, by showing how different sorts of emotion (primary, secondary and tertiary, to a first approximation) could arise from the functioning of different architectural layers in one individual, e.g. the evolutionarily oldest reactive mechanisms, newer deliberative mechanisms able to contemplate non-existent but possible states and processes, and meta-management mechanisms that monitor and control other internal processes. This is an oversimplification, but is mentioned as an illustration of the way in which conceptual analysis can be enriched if it is informed by theories about information-processing architectures. Our theory contradicts many of the popular views of emotions, e.g. theories treating emotions as involving bodily episodes and experienced mental episodes.

The ideas were further extended in collaboration with Brian Logan, then Matthias Scheutz and Ron Chrisley, e.g. in Sloman, Chrisley, and Scheutz (2005). In contrast with these ideas, the

[11] For more on this see this discussion of logical topography and logical geography:
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/logical-geography.html

vast majority of emotion researchers interested in computational models seem to treat emotions as definitionally linked to expressive behaviours, and actual mental episodes, rather than enduring but possibly dormant dispositions to produce such behaviours and episodes. I think it will eventually become clear, partly as a result of more sophisticated architecture-based analysis, that current work on affective states and processes focuses on only a tiny subset of the actual variety. Novelists, playwrights and poets are far more comprehensive in their grasp of the phenomena.

2 Examples

Question 2: What example(s) from your work (or the work of others) best illustrates the fruitful use of a computational and/or informational approach for foundational researches and/or applications?

It all depends on what 'fruitful' is supposed to mean. Not everyone will agree on what should be called fruitful. More importantly, it depends on what 'computational' and 'informational' are taken to mean. Neither of those has a clear and generally agreed meaning. As far as I am concerned what is fruitful is the approach that centres on the study of kinds of mechanisms that could explain *how systems work*, and *how processes occur*: the design-based approach. The processes in question may include physical and chemical processes, but are dominated by information-processing mechanisms, including probably kinds of mechanisms that we have not yet thought of.

2.1 My own work

Anyhow, almost all of my work since about 1971 has been based on the design stance, as illustrated above. Answering Question 2 would amount to reproducing most of the contents of my online account of topics I have worked on `http://www.cs.bham.ac.uk/~axs/my-doings.html`.

2.2 Work of others

Almost everything written by John McCarthy and Marvin Minsky, and most of Margaret Boden's work also illustrates this. Also much of the work by Herbert Simon and Allen Newell. A lot of the work by Andy Clark is relevant though as far as I know he has not engaged in developing designs for working systems: rather he works mainly at a methodological level. There are many more, of course, too many to list. In earlier times, even work by Immanuel Kant, Gottlob Frege, Jean Piaget, and Kenneth Craik, among others, was relevant, among others, though they lacked the conceptual tools we now have for thinking about information-processing systems, just as we now lack conceptual tools that will be developed in years to come.

3 The proper role?

Question 3: What is the proper role of computer science and/or information theory in relation to other disciplines, including other philosophical areas?

I would not use the phrase "proper role". I would say that there are a number of things about the science of information that are relevant to all aspects of science and philosophy. So people learning in detail about types of information-processing system may gain the following benefits.

- The most important single fact is that whereas much philosophical discussion has been about necessary or sufficient conditions for something to be the case, a great deal of such discussion has simply ignored the problem of explaining how *processes* can occur, and in particular explaining how mental mechanisms work: Above all else, minds *do things*, even if much of what they do is internal.

  Many people know that life cannot be understood if we ignore how its many forms evolved, how individual living things develop, how individuals interact with their environment, including other similar individuals. People simply take it for granted that life is about processes. But they fail to notice that minds *do things*, and that something that is always static and simply has information (e.g. beliefs, memories, concepts) and states that use information (e.g. desires, preferences, hopes, fears, intentions, etc.), without ever acquiring, manipulating, deriving, constructing and using information, could not be a mind. That means that understanding what minds are requires understanding how all those processes can happen. Unfortunately, when philosophers do try to discuss mechanisms that might explain such processes they often produce vague, incomplete, totally inadequate descriptions because they have no experience of the difference between using a specification to make something that works, and merely talking about such things: learning how to design, test, debug, analyse, compare and criticise working systems is not yet a standard part of the philosophical curriculum, though some of us attempted to produce a new breed of philosophy graduates educated in such matters at the university of Sussex, starting in the mid 1970s.

  So the first benefit is learning how to produce explanations of working systems (e.g. minds), in particular explanations that are capable of being tested in working implementations, as opposed to merely looking plausible to readers. This requires learning new ways to think about process and structure, i.e. how structures can produce and constrain processes and how processes can manipulate structures – including 3-D structures in the environment and abstract structures in information-processing systems.

- Another benefit is learning how to detect inadequate explanations of mental and biological phenomena, e.g. explanations that could not be used as the basis for designing or building a system of the kind whose abilities are being explained.

- Understanding the variety of types of virtual machine and the variety of ways in which virtual machines can be implemented or realised in physical machines or other virtual machines, will, I suspect, provide much matter for philosophical analysis in future years.

Those are some of the broad themes associated with learning about computer science, software engineering and the theory of information. More specific benefits include:

- Learning how to analyse designs for various kinds of virtual machine as a way of clarifying old philosophical questions, e.g. about types of freedom, the nature of mental states and processes, varieties of consciousness, etc.
- Understanding the deep implications of the differences between matter-manipulating, energy-manipulating and information-manipulating machines, and their relationships. E.g. information manipulating machines have to be implemented in energy manipulating machines. Matter and energy are conserved (or, to be more precise, mass-energy is conserved) and they are measurable. In contrast information is not conserved: I can give you much of my information without losing any of it. Moreover what is important about information is not usually a measurable quantity but its *content*, i.e. what is referred to, and what properties, relationships, generalisations, are expressed.
- Understanding the variety of *requirements* for machines (natural or artificial) with different combinations of information-processing competence produced by evolution, e.g. machines like microbes, like insects, like fishes, like reptiles, like birds, like mammals, of many kinds. This includes understanding the information-processing demands imposed both by different sorts of environment and by different sorts of physical design for the animal or machine, as explained in (Sloman, 2007)
- Understanding issues related to (decidability) solvability and complexity which are relevant to the differences between different mechanisms for doing the same task, and to differences between different tasks.
- Understanding the role of Turing machines and von Neumann machines within the larger space of possible information-processing machines, including the chemical information-processing machines used throughout living systems, and machines that grow their architectures.

It must be emphasised that not everyone who studies computing, software engineering or AI and who learns to write computer programs ends up understanding all the above issues. For example, there are now many AI researchers who, because of recent swings of fashion in the field, have no first-hand experience of designing any of the kinds of systems developed in symbolic AI, including systems that manipulate symbolic structures. As a result they know only about neural nets and mechanisms for operating on numbers, matrices, probability distributions and the like, especially mechanisms and processes that are closely coupled with sensori-motor systems. That is a good way to think about most information-processing in insects and other invertebrates, but completely ignores many of the interesting competences of birds, humans, and other mammals, including mathematical and philosophical competences of humans.

This leads to a very serious narrowing of vision regarding possible designs and the requirements to which they are suited. This restricted (and often prejudiced) understanding of information-processing systems can be found in researchers in several disciplines. I expect the publication of Boden (2006) will help to counter this narrowness, though it will take time.

4 Neglected topics

Question 4: What do you consider the most neglected topics and/or contributions in late 20th century studies of computation and/or information?

I have partially answered this above, e.g. in the last section. More generally, there are two sorts of task that have not been properly understood or pursued, with a resulting serious impact in our understanding of natural and artificial information-processing systems, namely:

1. understanding the variety of designs produced by evolution, including designs that make significant use of chemical information-processing, and designs that make significant use of interacting virtual machines;
2. understanding the requirements for specific designs, e.g. requirements for a human-like robot.

I think most philosophers ignore computers and computation except as providing tools to enable them to write papers, exchange email, browse the web etc. And they don't bother to learn how the machines they use are designed, implemented, debugged, and extended. One of the by-products of this neglect has been a failure of many philosophers to understand the variety of types of information-processing architecture and the way in which different architectures support different sets of concepts (Sloman, 2002a). In particular, as mentioned above, many of them have no understanding of the philosophical interest of relations between virtual machines and the underlying physical machines. A very recent paper by John Pollock (2008) is a rare exception. Dennett mentions virtual machines so often that he might be taken as an exception, but he makes it clear, e.g. in Dennett (2007) that he thinks that talk of virtual machines is a useful fiction, not to be taken seriously. In footnote 10 he asserts that it is standard practice in computer circles to allow virtual machine talk "to pass for the truth with no raised eyebrows, no need for the reminder that it is, officially, just metaphorical". This could not be further from the truth, as Pollock's paper confirms, as do many discussions among software engineers regarding bugs in complex systems. An event in a fictional machine could not cause an airliner to crash.

This limited understanding of the nature of complex virtual machines can lead at one extreme to gross over-simplification of the problems of explaining human and animal competences and designing machines with similar competences, and at the other extreme to a belief that the task of understanding humans and explaining how they work is far too complex and subtle to be aided by learning about information-processing, leading to varieties of "mysterianism" often inspired by arguments presented or reported in Chalmers (1996).

There have been many over-simplifications that result from inadequate understanding of requirements and designs, including the spread of "symbol grounding theory",[12] even among many people who work with computers, over-inflation of importance of emotions for intelligent systems,[13] including a failure to understand that emotions are just a subset among a large variety of types of control process, confusions about embodiment and situatedness, exaggerating the role of sensorimotor contingencies in cognition[14] and confusions about the relevance of Turing machines to the long term goals of research in AI (Sloman, 2002c).

5 Open problems

Question 5: What are the most important open problems concerning computation and/or information and what are the prospects for progress?

Importance is in the eye of the beholder. I can only talk about what interests me, and why.

I think the two hardest unsolved problems are (a) specifying what vision is and how human visual capabilities can be explained, and (b) analysing the concept of causation. I have already referred to some of the unsolved problems regarding vision. I conjecture that one of the reasons why it is so difficult to come up with a satisfactory analysis of causation is that there are two concepts of causation in wide use, namely Humean causation, and Kantian causation. Kantian causal reasoning is based on information about structures and composition of physical things, and often uses the kind of spatial reasoning that got me interested in AI. Some of the things philosophers do when discussing causation, such as talking about possible worlds, can divert attention from the task of understanding what happens when individuals *use* their understanding of causation in the process of learning things and achieving things, in the environment. At a workshop on natural and artificial cognition in June 2007 Jackie Chappell and I gave a linked pair of presentations on varieties of causation, available online at `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac`.

Other problems still in need of investigation, include: explaining how a "congenitally blind" human-like robot could "grow up", as congenitally blind humans do, achieving a rich understanding of the 3-D spatial world; and explaining what it is to be in various sorts of affective state, e.g. interested in something, liking or disliking something, valuing something, or finding something funny (one of the hardest cases). My original problem of explaining how understanding of mathematics, including the development through infancy and childhood of spatial forms of reasoning alongside reasoning using Fregean forms of representation, also remains unsolved. This includes explaining how a child or young robot can first discover a generalisation empirically, then come to understand that it is a *necessary* truth – for example that counting the same set in different orders must give the same result, or that a triangle and a circle can intersect in at most six points (Sloman, 1978a, 2008b) – an important step to justifying Kant's views on the nature of

[12] Criticised in this presentation `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models`

[13] As explained in this "popular" presentation `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cafe04`

[14] As explained in `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603`

mathematical knowledge. Closely connected with this is the problem of explaining how ontologies and forms of representation develop through playful exploration of both ideas and the environment (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0604).

There are more abstract problems about the diversity of sets of requirements (i.e. niches) and the diversity of designs and how feedback loops linking them can drive evolution. Understanding this fully (if that is possible) will require development of new mathematics. Other problems concern the sorts of architectures that are possible. It is very unlikely that we have an adequate characterisation of the space of information-processing architectures for intelligent systems, especially architectures that grow themselves.

There are many more unsolved problems and I am sure that researchers in this area will be kept busy for several more decades at least and possibly for centuries more. Confident pronouncements about how soon computer-based machines will overtake humans can safely be ignored while the problems that need to be solved are still not understood.

Perhaps the most important unsolved problem of all is how to undo the damage caused, and retrieve the opportunities lost, by all the practices in the last few decades perverting computing education in schools towards the goal of producing industry-fodder instead of attempting to stretch the creative, independent, model-building and reasoning capabilities of young minds.

Some related online papers and presentations

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703
Computational Cognitive Epigenetics

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0705
On qualia in intelligent machines and a new kind of Turing test.

http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cafe04
Do machines, natural or artificial, really need emotions?

http://www.cs.bham.ac.uk/~axs/talks.html
Online presentations since about 2000.

(I attempt to put all my publications on my web site, and will endeavour to ensure that they remain accessible there indefinitely. I wish more people would do the same. It is wrong that results of publicly funded research, freely provided, should not be freely available to anyone in the world.)

References

Boden, M. A. (1978). *Artificial intelligence and natural man.* Hassocks, Sussex: Harvester Press. (Second edition 1986. MIT Press)

Boden, M. A. (2006). *Mind As Machine: A history of Cognitive Science (Vols 1–2).* Oxford: Oxford University Pres.

Chalmers, D. J. (1996). *The conscious mind: In search of a fundamental theory*. New York, Oxford: Oxford University Press.
Chappell, J., & Sloman, A. (2007). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, *3*(3), 211–239. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609)
Dennett, D. C. (2007). Heterophenomenology reconsidered. *Phenomenology and the Cognitive Sciences*, *6*(1-2), 247–270. (DOI 10.1007/s11097-006-9044-9)
Gibson, E. J., & Pick, A. D. (2000). *An Ecological Approach to Perceptual Learning and Development*. New York: Oxford University Press.
Jablonka, E., & Lamb, M. J. (2005). *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge MA: MIT Press.
Jamnik, M., Bundy, A., & Green, I. (1999). On automating diagrammatic proofs of arithmetic arguments. *Journal of Logic, Language and Information*, *8*(3), 297–321.
Kant, I. (1781). *Critique of pure reason*. London: Macmillan. (Translated (1929) by Norman Kemp Smith)
McCarthy, J. (1996). *The Well Designed Child.* Stanford University. (http://www-formal.stanford.edu/jmc/child1.html)
McCarthy, J., & Hayes, P. (1969). Some philosophical problems from the standpoint of AI. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 4* (pp. 463–502). Edinburgh, Scotland: Edinburgh University Press. (http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html)
Minsky, M. L. (2006). *The Emotion Machine*. New York: Pantheon.
Pollock, J. L. (2008). What Am I? Virtual machines and the mind/body problem. *Philosophy and Phenomenological Research.*, *76*(2), 237–309. (http://philsci-archive.pitt.edu/archive/00003341)
Sloman, A. (1962). *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth*. Unpublished doctoral dissertation, Oxford University. (http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#706)
Sloman, A. (1965a). Functions and Rogators. In J. N. Crossley & M. A. E. Dummett (Eds.), *Formal Systems and Recursive Functions: Proceedings of the Eighth Logic Colloquium Oxford, July 1963* (pp. 156–175). Amsterdaim: North-Holland Publishing Co.
Sloman, A. (1965b). 'Necessary', 'A Priori' and 'Analytic'. *Analysis*, *26*(1), 12–16. (Now online http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#701)
Sloman, A. (1968/9.). Explaining Logical Necessity. *Proceedings of the Aristotelian Society*, *69*, 33-50. (http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#712)
Sloman, A. (1971). Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In *Proc 2nd ijcai* (pp. 209–226). London: William Kaufmann. (http://www.cs.bham.ac.uk/research/cogaff/04.html#200407)
Sloman, A. (1978a). *The computer revolution in philosophy*. Hassocks, Sussex: Harvester Press (and Humanities Press). (http://www.cs.bham.ac.uk/research/cogaff/crp)
Sloman, A. (1978b). What About Their Internal Languages? Commentary on three articles by Premack, D., Woodruff, G., by Griffin, D.R., and by Savage-Rumbaugh, E.S., Rumbaugh,

D.R., Boysen, S. in *Behavioral and Brain Sciences* Journal 1978, 1 (4). *Behavioral and Brain Sciences*, *1*(4), 515. (http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#713)

Sloman, A. (1979). The primacy of non-communicative language. In M. MacCafferty & K. Gray (Eds.), *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979* (pp. 1–15). London: Aslib. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43)

Sloman, A. (1982). Image interpretation: The way ahead? In O. Braddick & A. Sleigh. (Eds.), *Physical and Biological Processing of Images (Proceedings of an international symposium organised by The Rank Prize Funds, London, 1982.)* (pp. 380–401). Berlin: Springer-Verlag. (http://www.cs.bham.ac.uk/research/projects/cogaff/06.html#0604)

Sloman, A. (1984). Experiencing Computation: A Tribute to Max Clowes. In M. Yazdani (Ed.), *New horizons in educational computing* (pp. 207 – 219). Chichester: Ellis Horwood Series In Artificial Intelligence. (http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#71)

Sloman, A. (1989). On designing a visual system (towards a gibsonian computational model of vision). *Journal of Experimental and Theoretical AI*, *1*(4), 289–337. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#7)

Sloman, A. (1993). The mind as a control system. In C. Hookway & D. Peterson (Eds.), *Philosophy and the cognitive sciences* (pp. 69–110). Cambridge, UK: Cambridge University Press. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#18)

Sloman, A. (1994). How to design a visual system – Gibson remembered. In D.Vernon (Ed.), *Computer vision: Craft, engineering and science.* Berlin: Springer Verlag.

Sloman, A. (1995). Musings on the roles of logical and non-logical representations in intelligence. In J. Glasgow, H. Narayanan, & B. Chandrasekaran (Eds.), *Diagrammatic reasoning: Computational and cognitive perspectives* (pp. 7–33). MIT Press.

Sloman, A. (1996). Towards a general theory of representations. In D.M.Peterson (Ed.), *Forms of representation: an interdisciplinary theme for cognitive science* (pp. 118–140). Exeter, U.K.: Intellect Books.

Sloman, A. (1998, August). Diagrams in the mind. In *in proceedings twd98 (thinking with diagrams: Is there a science of diagrams?)* (pp. 1–9). Aberystwyth.

Sloman, A. (2001). Evolvable biologically plausible visual architectures. In T. Cootes & C. Taylor (Eds.), *Proceedings of British Machine Vision Conference* (pp. 313–322). Manchester: BMVA.

Sloman, A. (2002a). Architecture-based conceptions of mind. In *In the Scope of Logic, Methodology, and Philosophy of Science (Vol II)* (pp. 403–427). Dordrecht: Kluwer. (http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#57)

Sloman, A. (2002b). Diagrams in the mind. In M. Anderson, B. Meyer, & P. Olivier (Eds.), *Diagrammatic representation and reasoning.* Berlin: Springer-Verlag.

Sloman, A. (2002c). The irrelevance of Turing machines to AI. In M. Scheutz (Ed.), *Computationalism: New Directions* (pp. 87–127). Cambridge, MA: MIT Press. (http://www.cs.bham.ac.uk/research/cogaff/00-02.html#77)

Sloman, A. (2006, May). *Sensorimotor vs objective contingencies* (Research Note No.

COSY-DP-0603). Birmingham, UK: School of Computer Science, University of Birmingham. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603)

Sloman, A. (2007). Diversity of Developmental Trajectories in Natural and Artificial Intelligence. In C. T. Morrison & T. T. Oates (Eds.), *Computational Approaches to Representation Change during Learning and Development. AAAI Fall Symposium 2007, Technical Report FS-07-03* (pp. 70–79). Menlo Park, CA: AAAI Press. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0704)

Sloman, A. (2008a, March). *Architectural and representational requirements for seeing processes, proto-affordances and affordances* (Research paper No. COSY-TR-0801a). UK: School of Computer Science, University of Birmingham. (Also presented at Dagstuhl workshop on *Logic and Probability for Scene Interpretation*, http://www.cs.bham.ac.uk/research/projects/cosy/papers#tr0801a)

Sloman, A. (2008b, March). *Kantian Philosophy of Mathematics and Young Robots* (To appear in proceedings MKM08 No. COSY-TR-0802). UK: School of Computer Science, University of Birmingham. (http://www.cs.bham.ac.uk/research/projects/cosy/papers#tr0802)

Sloman, A., & Chappell, J. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings IJCAI'05* (pp. 1187–1192). Edinburgh: IJCAI. (http://www.cs.bham.ac.uk/research/cogaff/05.html#200502)

Sloman, A., & Chappell, J. (2007). Computational Cognitive Epigenetics (Commentary on (Jablonka & Lamb, 2005)). *Behavioral and Brain Sciences*, *30*(4), 375–6. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703)

Sloman, A., Chrisley, R., & Scheutz, M. (2005). The architectural basis of affective states and processes. In M. Arbib & J.-M. Fellous (Eds.), *Who Needs Emotions?: The Brain Meets the Robot* (pp. 203–244). New York: Oxford University Press. (http://www.cs.bham.ac.uk/research/cogaff/03.html#200305)

Trehub, A. (1991). *The Cognitive Brain*. Cambridge, MA: MIT Press. (http://www.people.umass.edu/trehub/)

Wegner, P. (1968). *Programming Languages, Information Structures, and Machine Organization*. New York: McGraw Hill.

Wing, J. M. (2006). Computational Thinking. *CACM*, *49*(3), 33–35. (http://www.cs.cmu.edu/afs/cs/usr/wing/www/publications/Wing06.pdf)

Winterstein, D. (2005). *Using Diagrammatic Reasoning for Theorem Proving in a Continuous Domain*. Unpublished doctoral dissertation, University of Edinburgh, School of Informatics. (http://www.era.lib.ed.ac.uk/handle/1842/642)

Wright, I., Sloman, A., & Beaudoin, L. (1996). Towards a design-based analysis of emotional episodes. *Philosophy Psychiatry and Psychology*, *3*(2), 101–126. (http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#2)