

http://www.bcs.org/sg/ccs/events_se.html

11th October 2002

(Slightly re-formatted 10 Dec 2016)

Science Museum, London

Artificial Intelligence Recollections of the pioneers

Early British AI AI and the study of mind

Aaron Sloman

School of Computer Science,
The University of Birmingham, UK

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

THIS PDF PRESENTATION IS ON-LINE AS "TALK 17" AT

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

(Last changed December 11, 2016)

ACKNOWLEDGEMENT

This talk was prepared using free software,
Linux, Latex, Acroread, Poplog, etc.

Diagrams were created using tgif, freely available from

`http://bourbon.cs.umd.edu:8001/tgif`

I am especially grateful to the developers of Linux
and all the gnu tools used in Linux

The CCS Meeting

The meeting at which this talk was presented was held on Friday 11th October at the Science Museum London.

It was organised by the Computer Conservation Society, a co-operative venture between the British Computer Society, the Science Museum of London, and the Museum of Science and Industry in Manchester.

<http://www.bcs.org/sg/ccs/>

It is dedicated to the conservation and restoration of early computers, preserving early computer software and other digital records of historical importance, and to recording the history of computing. It encourages research into, and awareness of, the importance of historic computers and early computer developments in Britain.

The meeting started with a video presentation by Donald Michie, who is in Australia, followed by talks by Jim Doran (Essex University), Austin Tate (Edinburgh University) and Aaron Sloman (Birmingham University).

Details at

http://www.bcs.org/sg/ccs/events_se.html

At the meeting I promised to expand my slides and make them generally available. This is the result. I have added five slides with autobiographical notes and a slide on Max Clowes. See also

<http://www.cs.bham.ac.uk/research/cogaff/sloman-clowestribute.html>

My entry into AI

My first degree (Cape Town, 1956) was in mathematics and physics. I went to Oxford to continue studying mathematics, but did not like some of the compulsory courses (fluid dynamics and differential equations) and switched to a D.Phil in mathematical logic, supervised by Hao Wang, whom I met rarely because he was doing a lot of work for IBM in the USA, developing theorem proving techniques.

I was gradually seduced by philosophy, and eventually switched to a D.Phil in philosophy (defending some of Kant's ideas about mathematical knowledge), completed in 1962. I was interviewed on the same day in 1962 for a job at IBM and a job in Hull University. The IBM interviewer told me I was better suited for a university, and for two years I taught logic and philosophy of science at Hull, then went to Sussex in 1964 where I taught a variety of philosophy courses and did research on meaning, reasoning, ethics (the logic of 'better') and the nature of mind.

Somehow I met Max Clowes, a brilliant and charismatic AI researcher in vision, around 1969, when he was as a research fellow in the Experimental Psychology Laboratory. Max, Keith Oatley (a psychologist now in Toronto) and I had regular meetings discussing how we were going to revolutionise the study of mind. I sat in on Max's undergraduate AI courses, learnt some Fortran and Algol programming (the only options available then) and was constantly harangued by Max about the need to drop my way of doing philosophy and instead take up AI programming.

My entry ...continued (2)

Max bullied me into submitting my first AI publication to the second IJCAI, in London 1971 (a paper criticising the logicist approach to AI), as a result of which Bernard Meltzer got funds from the Science Research Council to bring me, with my family, to Edinburgh for a wonderful year, in 1972-3.

I met many excellent researchers as I roamed around the various departments (Dept. of Machine Intelligence and Perception, led by Donald Michie, Bionics, led by Jim Howe, the Epistemics group led by Christopher Longuet-Higgins and the Department of Computational Logic, led by Bernard Meltzer, my very generous host.) The fact that there were many political tensions between the leaders did not affect communication between junior researchers and visitors.

I met many PhD students who became famous researchers (including Geoff Hinton, Alan Bundy, Austin Tate, Richard Power, Gordon Plotkin) and they and the other researchers and visitors taught me a huge amount. Bob Boyer and J Moore were working on their subsequently famous Lisp theorem prover: they gave me a five hour tutorial on Lisp, the most intensive learning experience I can remember. Julian Davis let me read the Pop2 code for Popler, his implementation of Carl Hewitt's "Planner" system. I learnt a huge amount about programming, and the power of Pop2 as a language, from Julian. He later went to Canada, and alas died there. Bob Kowalski introduced me to logic programming and connection graphs. Sylvia Weir (who later went to MIT) became a friend and general guide and tutor.

My entry ...continued (3)

Many excellent researchers visited Edinburgh that year, including Danny Bobrow who taught me Logo before going off to help found Xerox PARC, Chris Brown (who worked on Freddy then went to Rochester), Woody Bledsoe (visiting from Texas), Alan Robinson (whose ideas led to the birth of Prolog), Seymour Papert (visiting from MIT), and others. I learnt more in that year in Edinburgh than in any year since about the age of four.

Thanks to pioneering work by Steve Salter, Robin Popplestone, Pat Ambler, Harry Barrow, Chris Brown, Rod Burstall and others, Freddy, the Edinburgh robot was taught to assemble a toy wooden car, refuting in advance many claims made later on about the narrow assumptions of early AI researchers.

In particular, the critics have no idea what it was like to use such inadequate computing facilities. It took about 10 minutes or more merely to find edges and regions in an image in order to recognize a cup and saucer. The only acceptable form of demo was a speeded up film! Concurrent visual perception and action was out of the question until many years later.

I went back to Sussex in October 1973 and there, with Maggie Boden (who had been writing about AI without my knowing it, and later published *Artificial Intelligence and Natural Man*), Max Clowes and Alistair Chalmers, a psychologist, founded an undergraduate teaching programme in Cognitive Studies, which later grew into the School of Cognitive and Computing Sciences (COGS).

My entry ...continued (4)

The main teaching and research language used for AI at Sussex for many years was Pop2, in various dialects and implementations. Initially it was accessed by researchers on a PDP-10 in Edinburgh via modem links, and by undergraduates on the local ICL 1900 machine running the George 3 operating system, using a ridiculous batch submission system. On the PDP-10 we first used Pop-10, implemented by Julian Davies, then WPOP (WonderPop) implemented by Robert Rae, later helped by Allan Ramsay (now at UMIST).

Meanwhile Steve Hardy wrote Pop-11, a reduced version, of Pop2, expanded with a pattern matcher which made (and still makes) a huge difference for AI teaching and research. Pop-11 was implemented in PDP11 assembler and ran on a DEC PDP11/40 computer from about 1976, at first under RX11-D, soon abandoned for Unix. Several universities took copies of Unix Pop-11 and used it for teaching and research. Around 1980/81 it was subsequently ported to Vax/VMS, with a full incremental compiler for Pop-11, instead of an interpreter for intermediate code.

Later, around 1982, an incremental compiler for Prolog was added by Chris Mellish, then developed by John Gibson and others). We called the combination Poplog. In 1983, The Japanese 5th Generation computer project and the UK Alvey programme generated much interest in AI, and since there were no good AI tools available on a Vax, interest in Poplog grew, and it actually funded a great deal of the growth of AI and Cognitive Sciences at Sussex.

My entry ...continued (5)

There's a long story about how marketing of Poplog was taken over by Systems Designers Ltd, around 1983, then by a spin-off company Integral Solutions Ltd in 1989 founded by Alan Montgomery and colleagues from SD. On the way, incremental compilers for Common Lisp and Standard ML were added, then later powerful graphical tools, and Poplog became one of the most versatile and sophisticated AI development environments in existence. Clementine, now owned by SPSS is the most successful commercial product based on Poplog, but there were others. I have a plaque awarded to ISL and Sussex by SMART for achieving 5 million dollars worth of sales of Poplog by the early 1990s.

I managed the poplog development team up to 1991, when I left Sussex for Birmingham, though most of the hard work was done by others (especially John Gibson, John Cunningham, Chris Mellish, John Williams, Roger Evans, Robert Duncan, Simon Nichols, Robert Smith, Ian Rogers, Mark and Ben Rubinstein, Tom Khabaza)

Since moving to Birmingham, I have continued developing Pop-11 based tools to support teaching and research on agent architectures (the SimAgent toolkit). These tools were and are used in Birmingham, at DERA, and some universities in the UK and USA. For more on all this see

<http://www.cs.bham.ac.uk/research/cogaff/sussex-pop11.html>

<http://www.cs.bham.ac.uk/research/poplog/freepoplog.html>

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

Max Clowes

Max was brilliant but impatient. He left Sussex around 1980, feeling that progress in AI was too slow. He likened it to inching oneself up a vertical ice-face.

He decided instead to try to bring AI into school teaching, and joined a teacher-training college in Devon.

Unfortunately he died soon after of a heart attack in 1981, the year another great vision researcher died, David Marr.

The editor of the the journal *Computing in Schools* asked me to write a tribute to Max, which was published in 1981.

An extended version is now available online here, in two formats:

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

<http://www.cs.bham.ac.uk/research/cogaff/sloman-clowestribute.html>

This tribute was much enlarged in March and April 2014, including

- recollections by Wendy Manktellow (nee Taylor), and
- a draft annotated biography of Max Clowes with publications.

END OF DIGRESSION:

The remaining slides constitute the presentation at the CCS meeting.

Disclaimer: I am not one of the pioneers!

AI had already been going for some time by the time I got involved, around 1969/70, under the influence of Max Clowes.

- Grey Walter's tortoise was well known (is that AI?)
- Minsky had written 'Steps towards artificial intelligence' a decade earlier.
- The General Problem Solver of Newell, Shaw and Simon was several years old.
- Selfridge's Pandemonium was well known.
- The book *Computers and Thought* had been published (1963)
- The Stanford Robot had done some trundling
<http://www.frc.ri.cmu.edu/~hpm/talks/robot.evolution.html>
- Simon had written his paper on motivation and emotions
- Checkers (draughts) programs were already beating humans.
- Various image interpretation programs had been developed
- Lisp was already a major AI language
- Pop2 was born and growing (later versions included Pop-10, Wpop, Pop-11)
- But:

By then computers were still pathetically inadequate for most AI tasks in speed, physical memory, address space, etc. and cost a fortune.

All this did not hinder my own attempts to expand deep interactions between AI and philosophy (1978).

Things were different in 1970

- AI was dominated by logic (almost all the theoretical work).

Example

J. McCarthy & P.J. Hayes, (1969,)

Some philosophical problems from the standpoint of AI,
in *Machine Intelligence 4*

- My own first AI publication in 1971 was an attack on that position, presented at the 2nd International Joint Conference on AI, at Imperial College London.

<http://www.cs.bham.ac.uk/research/cogaff/62-80.html#1971-02>

E.g. I argued that diagrammatic reasoning was different from and sometimes more useful than logical reasoning.

Others have made the same point, before and since then.

- I've been arguing with McCarthy and Hayes ever since, though it is now widely accepted that we need a wide class of forms of representation to do AI.

But progress has been much slower than everyone hoped.

WHY???

First we have to understand the goals of AI.

WHAT IS ARTIFICIAL INTELLIGENCE?

?????

?????

WHAT IS ARTIFICIAL INTELLIGENCE?

**Getting machines to behave
like machines in the movies???**

Including having emotions???

WHAT IS ARTIFICIAL INTELLIGENCE?

A better answer:

AI is a (relatively) new approach to some very old questions:

WHAT ARE MINDS AND HOW ARE THEY RELATED TO BODIES?

WHAT IS INTELLIGENCE, AND CAN MACHINES HAVE IT?

WHAT IS CONSCIOUSNESS?

WHAT IS ARTIFICIAL INTELLIGENCE?

A better answer:

AI is a (relatively) new approach to some very old questions:

WHAT ARE MINDS AND HOW ARE THEY RELATED TO BODIES?

WHAT IS INTELLIGENCE, AND CAN MACHINES HAVE IT?

WHAT IS CONSCIOUSNESS?

It combines with and contributes to several other disciplines, including:

- psychology,
- philosophy,
- linguistics,
- biology,
- anthropology,
- logic,
- mathematics,
- computer science
- software engineering,
- mechanical engineering,

Compare some older views of mind:

- A ghost in a machine (dualism)
- Some of them just ghosts, without machines ?
- Social/political models of mind (Plato, Freud)
- Mechanical models (e.g. levers, steam engines)
- Electrical models (old telephone exchanges)

Compare some older views of mind:

- A ghost in a machine (dualism)
- Some of them just ghosts, without machines ?
- Social/political models of mind (Plato)
- Mechanical models (e.g. levers, steam engines)
- Electrical models (old telephone exchanges)

PROBLEMS WITH OLDER APPROACHES

- Some lack explanatory power (ghost in the machine)
 - **Any intelligent ghost must contain a machine**
- Some are circular (Social/Political models of mind)
- Some offer explanations that are too crude to explain fine detail and do not generalise (e.g. mechanical and electrical models)

Compare some older views of mind:

- A ghost in a machine (dualism)
- Some of them just ghosts, without machines ?
- Social/political models of mind (Plato)
- Mechanical models (e.g. levers, steam engines)
- Electrical models (old telephone exchanges)

PROBLEMS WITH OLDER APPROACHES

- Some lack explanatory power (ghost in the machine)
 - **Any intelligent ghost must contain a machine**
- Some are circular (Social/Political models of mind)
- Some offer explanations that are too crude to explain fine detail and do not generalise (e.g. mechanical and electrical models)

Only during the last century did we begin to find good ways to think about and model systems that process information, especially information with rich implications, some of which is complex and rapidly changing.

WE STILL HAVE MUCH TO LEARN.

Some History

From the earliest days there were two sorts of motivations in AI

- **engineering motivations** – making new useful things
- **scientific motivations** – explaining what already exists and could exist

Some History

From the earliest days there were two sorts of motivations in AI

- **engineering motivations** – making new useful things
- **scientific motivations** – explaining what already exists and could exist

Alan Turing was interested in both

as shown in his 1950 paper *Computing machinery and intelligence*

now available online at <http://www.abelard.org/turpap/turpap.htm>

The other main founders of AI were all more interested in AI as science than AI as engineering. For example:

- John McCarthy
- Marvin Minsky
- Herbert Simon (Nobel prize winner)

Minsky and (especially) Simon in particular were interested in AI as a new paradigm for studying human minds.

Even McCarthy starts from observations about human minds.

AI is not a branch of Computer Science

AI considered as science is different from Computer Science.

- CS studies forms of computation **in general** and their properties.
- AI also studies and attempts to explain **natural** information processing systems.
 - E.g. humans and other animals. Maybe martians, etc.?
- AI uses CS
- And feeds new problems, concepts and techniques into CS.
- That's something like the way in which physics relates to mathematics.
- But physics is not a branch of mathematics, even though there is no clear division between applied mathematics and physics.

AI overlaps with several disciplines and is not just a part of any of them.

False hopes

Early AI researchers were grossly over-optimistic.

They made predictions that did not come true:
a failing of several of the publicists for AI.

WHY?

Because

- People did not understand the difficulty, complexity and variety of the tasks.
- They did not know enough philosophy, linguistics, psychology, biology...
- It was too easy to assume that we already **know** what humans can do, leaving only the task of modelling and explaining human capabilities.
- Ideas about software architectures were still too primitive.
- And

False hopes

Early AI researchers were grossly over-optimistic.

They made predictions that did not come true:
a failing of several of the publicists for AI.

WHY?

Because

- People did not understand the difficulty, complexity and variety of the tasks.
- They did not know enough philosophy, linguistics, psychology, biology...
- It was too easy to assume that we already **know** what humans can do, leaving only the task of modelling and explaining human capabilities.
- Ideas about software architectures were still too primitive.
- Computers were pathetically inadequate for many years (except for limited classes of tasks).

It took Freddy the Edinburgh robot at least 10 minutes to find the edges in an image of a cup on a saucer, around 1972. See

<http://www.ipab.informatics.ed.ac.uk/IAS.html>

<http://www-robotics.cs.umass.edu/~pop/VAP.html>

But there was progress: we were learning

Despite the slow progress and disappointments, early results of AI research helped to reveal previously unnoticed subtleties of human (and animal) intelligence, many of which come to light as a result of finding that computer models are un-human-like in unexpected ways.

E.g. see

Margaret Boden (1977) *Artificial Intelligence and Natural Man*

(my 1978 book: *The computer revolution in philosophy* now online free.

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

Around 1971 Terry Winograd's PhD thesis was published.

DEMO OF SIMPLIFIED SHRDLU AVAILABLE (USING POP-11).

Interplay of syntax, semantics, world knowledge, planning...

(This illustrates deliberative mechanisms, mentioned below.)

There have been many small successes in AI

Including

- many applications in a wide variety of software systems
- in robotics
- most recently applications of AI in computer entertainments, e.g. modelling emotional processes.

Conjecture:

Computer entertainments involving synthetic characters or robots will be a major driving force for AI in the near future, partly because the problems are very challenging and interesting, partly because the revenue prospects run into billions of dollars.

BUT A GREAT DEAL OF VERY DIFFICULT WORK REMAINS TO BE DONE.

AI vision systems, for instance, have achieved very little of what human vision can do, or even the visual capabilities of squirrels, monkeys and nest-building birds.

In part the problem is understanding what human and animal visual systems can do, for instance perceiving more than just physical objects properties and relationships. (Compare Gibson on “affordances”.)

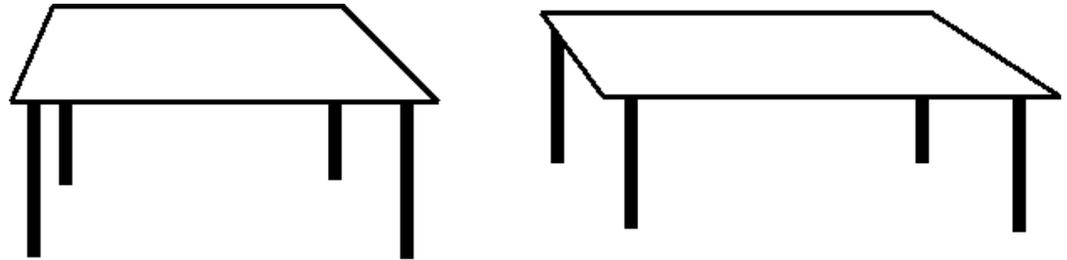
Added 2016: “Perception of impossibilities”:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

Examples of research problems in AI

Vision – perhaps the hardest problem in AI

How do we get from 2-D patterns of illumination on our retinas to percepts of a 3-D world?



What if the objects are flexible, irregularly shaped, and moving?

How do we see expressions of emotion in faces?

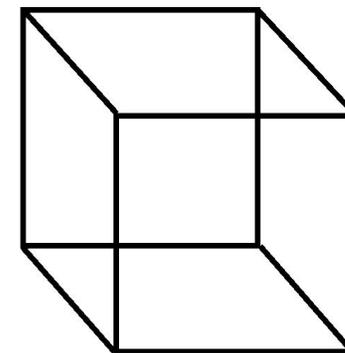
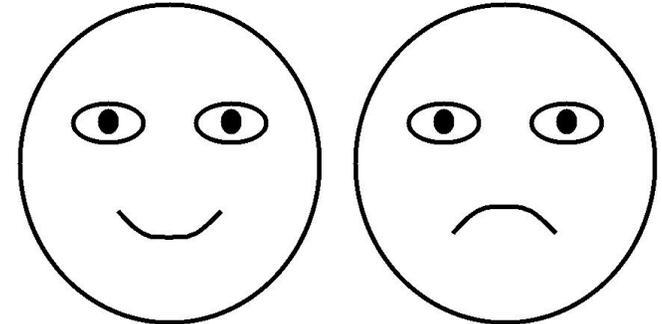
How are emotions represented in perceivers?

Why do the eyes look different if you gaze at them?

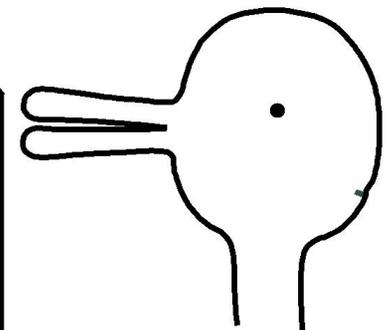
How can we see the same 2-D visual input in different ways?

How are the differences represented?

And many more, including perception of motion, visual pleasure, mathematical discovery, study of motivation and emotions, etc.



Necker cube



Duck-rabbit

Oct 2002

There are many other AI research areas

How does human language work?

- *How do we understand sentences we have never heard before?*
- *How do we produce sentences we have never heard before?*
- *How can we think about and talk about things that do not exist (e.g. Father Christmas, Harry Potter, Darth Vader, heaven?)*

How do we do mathematical reasoning

- *including thinking about infinite sets (like the set 2, 4, 6, 8, 10, 12,)?*

How do we select and control our actions?

How do we make plans

How do we learn (concepts, theories, skills, languages, ...)

Some of these questions arise for other animals also.
And also for robots and intelligent software systems.

AI tools and techniques

There have been many swings of fashion and spurious battles between factions.

Ignoring silly squabbles about “the one right way” we can say that AI researchers have explored many different forms of information processing (different forms of representation, different forms of searching, different ways of using knowledge, different types of learning), e.g.

- Symbol-manipulating programs that build structural descriptions
(SHRDLU demo)
- Logical reasoning systems and other deliberative systems
- Neural nets and other stochastic mechanisms
- Behaviour based (non deliberative) systems
(Sheepdog demo)
- Genetic algorithms (evolutionary computation)
- Dynamical systems (based on models from physics)
(Braitenberg demo — not presented at the meeting)
- Where appropriate, new sorts of hardware

Do not believe anyone who claims that only one kind of technique or mechanism works: some ignorant people have narrow-minded views of the nature and diversity of the problems.

AI languages

To support the development of these and other techniques, AI researchers have designed especially powerful and flexible programming languages, e.g.

SAIL (Stanford AI Language, now defunct)
Lisp, Prolog, Pop2, Pop-11, Scheme, T,
Rule-based languages, constraint languages, ...

The AI languages differ in important ways from more widely used languages such as C, C++, Java, ...

However, people have begun to implement AI toolkits on top of the other languages (e.g. JESS, Jack).

For more on the special requirements for AI languages see:

<http://www.cs.bham.ac.uk/~axs/misc/talks/#talk11>
<http://www.cs.bham.ac.uk/research/poplog/primer/>

It is likely that as the needs are more widely appreciated, key features of AI languages will be added to other popular languages, e.g. Java (which already has one major feature, garbage collection, once regarded by people outside AI as a luxury for the idle rich).

POPLOG: A UK product

A particularly important system was Poplog, a multi-language system combining ideas from Pop2 (now Pop-11) Prolog, Lisp and Standard ML. The Alvey program and the SERC helped to fund its development. Some of its history is described here:

<http://www.cs.bham.ac.uk/research/cogaff/sussex-pop11.html>

It was used by ISL for the development of Clementine, the well known data-mining system (eventually bought by SPSS).

Poplog used to be a very expensive system, which hindered its take-up. It has recently been made available as free, open source software:

<http://www.cs.bham.ac.uk/research/poplog/freepoplog.html>

Because there is a linux version it is starting to take off again.

There is also a windows version, but no graphics for it yet.

Poplog may be further developed if we can find funds — or even if we can't.

Some recent developments

In the last decade and a half, there has been a growing realisation of the need to understand **Architectures** in which various kinds structures and mechanisms are combined to form complex wholes with many cooperating parts

Examples

- SOAR
- PRS (For agents with beliefs, desires, intentions)
- JACK (distributed version of PRS)
- ACT/RP
- Minsky's Society of Mind and The Emotion machine
see his web site
- Swarm-like systems
- Our H-CogAff architecture
- and many others, including many distributed architectures

Towards the future

I believe that we now need to think hard about
[The space of possible architectures.](#)

For this we have developed an architecture-schema CogAff, which subsumes many different architectures, that can be explored so as to understand the trade-offs.

See our papers

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

and online presentations

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

Our multi-paradigm SimAgent toolkit aids the exploration of diverse architectures. See

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

Within the space of possibilities generated by the CogAff schema we have been trying to devise an architecture to encompass a wide range of human capabilities: H-CogAff

H-Cogaff – a first draft model of your mind

The Birmingham “CogAff” project has been developing a framework for characterising a wide variety of types of minds, of humans, other animals, and possible future robots.

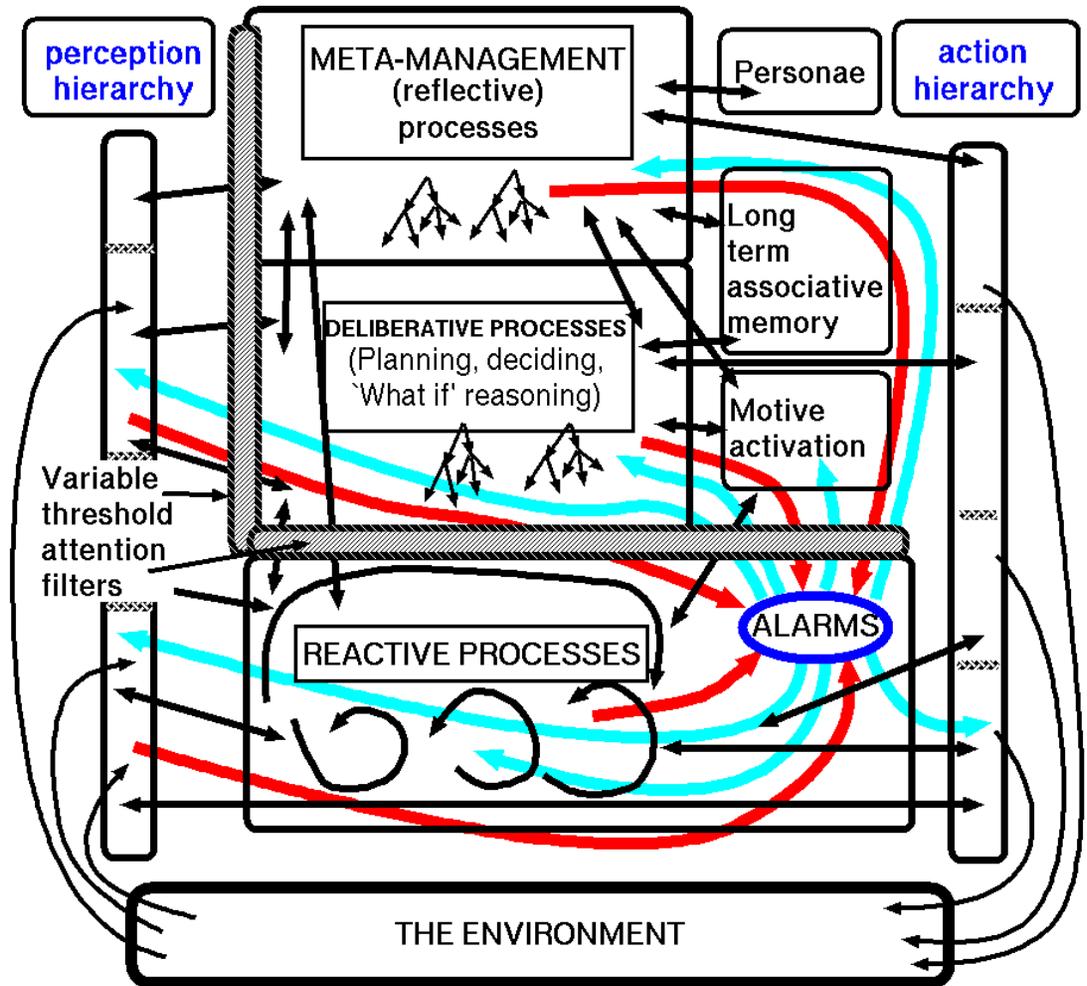
The framework incorporates evolutionarily ancient mechanisms co-existing and co-operating or competing with new mechanisms capable of doing different tasks (e.g. reasoning about what might happen). The figure gives an “impressionistic” overview of some of the complexity in our first draft H-CogAff architecture.

E.g. different sorts of emotions are generated in different levels.

More details including papers, slide presentations and software tools can be found at

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

CCS



Misplaced criticism

There there are critics who think the task of AI is beyond the scope of computer-based models, for instance because they think that animal brains are essentially different from computers.

Of course they are different. But whether that difference is enough to rule out replication of every possible variety of human thoughts, percepts, memories, feelings, emotions is not clear. That's because

- We do not yet know enough about what minds and brains are.
- We don't yet have a full grasp of what computers can and cannot do, especially when we connect many of them together and let them run asynchronously.

So we can't settle the questions yet. Sometimes people try to settle the question without waiting to find out the truth because they are *afraid* of certain answers.

These topics are discussed at greater length in a presentation on "Architectures and the spaces they inhabit" to two seminars at IBM in March 2002, and presentations in my talks directory:

<http://www.cs.bham.ac.uk/research/cogaff/ibm02>

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

An invitation

There remain many very difficult and fascinating scientific and engineering problems.

Why not join the fun?

Note added 11 Dec 2016

Since late 2011 my own work on AI expanded (under the influence of Turing's 1952 paper on Morphogenesis) to include **The Meta-Morphogenesis Project**

<http://goo.gl/9eN8Ks>