Slides presented at
KI2006 Symposium: 50 years of AI
17 Jun 2006
(Expanded after the conference)

Fundamental Questions - The second decade

Towards Architectures for Human-like Machines

Aaron Sloman

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

School of Computer Science, The University of Birmingham, UK

Online here:
http://www.cs.bham.ac.uk/research/cogaff/talks/#ki2006

Other related presentations:
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/

Papers and tools:
http://www.cs.bham.ac.uk/research/projects/cogaff/

A view from the side

Thanks to Max Clowes, my education in AI started in 1969 when I was a young lecturer in philosophy at the University of Sussex

I started to learn to program, attended Max's lectures and began reading many things – and was especially impressed by

Marvin Minsky's long paper

'Steps towards Artificial Intelligence',

published 1963 in the book *Computers and Thought*, eds. Feigenbaum and Feldman,

and also a collection of papers he edited in 1968

Semantic Information Processing

both of which should still be compulsory reading for everyone interested in natural or artificial minds.

INTERESTING WORK OF EARLY 70s

Around the time I started learning AI

Sussman's HACKER

Patrick Winston produced a thesis on inducing structural descriptions from examples (extending work by TG Evans mentioned in Minsky's presentation).

Winograd produced a thesis on understanding natural language by interleaving use of syntax, semantics and world knowledge.

Show simplified SHRDLU demo (Available in Poplog).

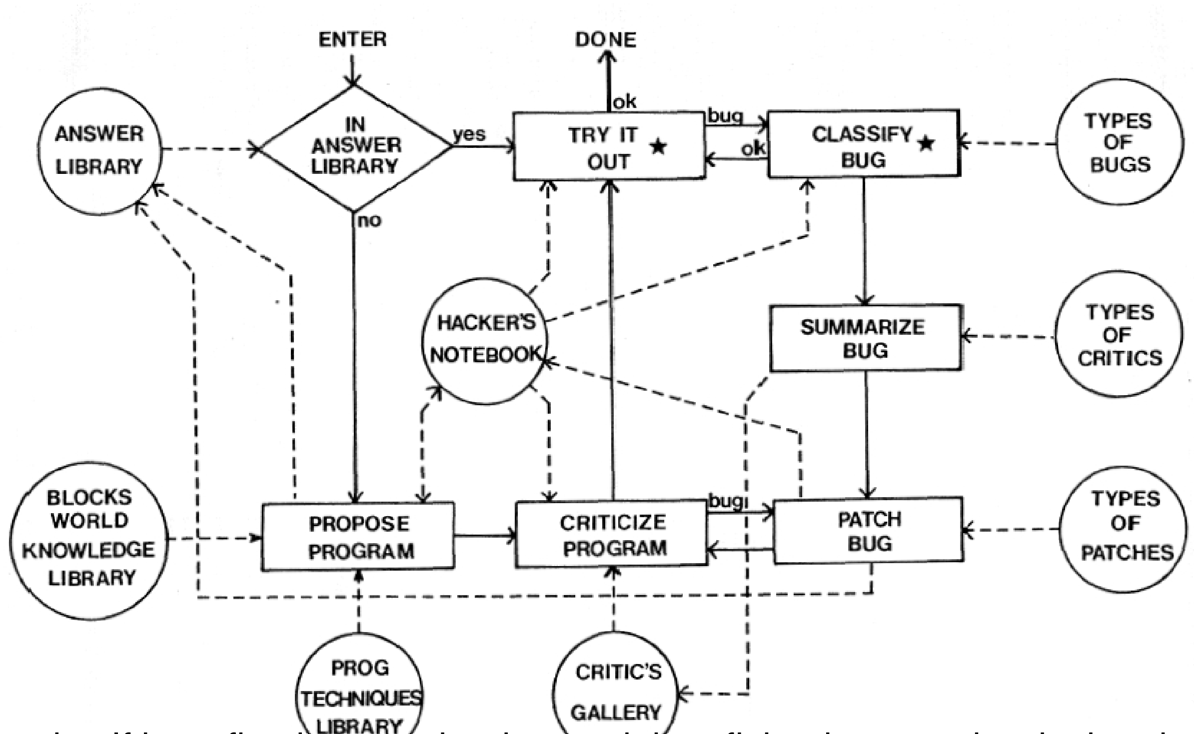

Sussman produced a PhD thesis reporting on HACKER, a design for a planning system that could debug itself by reflecting on what it was doing, fixing bugs and reducing the need to debug in future.
(Thesis now online at MIT, also a book *A computational model of skill acquisition* 1975, Elsevier.)

Show Josh yogurt video: an 11 month old feeds its mind while feeding its belly. Feeding belly uses spoon, yogurt, mouth and hands. Feeding mind uses all those and also legs, fingers, carpet.

`http://www.cs.bham.ac.uk/~axs/fig/josh23_0040.mpg`

AI and Philosophy

Within a couple of years I realised that the best way to do philosophy was to do AI:

I.e. design and implement fragments of working minds in order to test out philosophical theories

- about meaning,
- about knowledge,
- about explanation,
- about the nature of science,
- about the nature of mathematics,
- about the nature of mind,
- about the mind-body relationship,
- about aesthetics,
- and many other things.

All of these represent complex interactions between structures and processes, some in physical machines same in virtual machines.

Previously philosophers discussed ‘necessary’ or ‘sufficient’ conditions for things, rarely how things worked, so as to satisfy those, or other conditions.

Kant tried to discuss how things work, but did not have the right conceptual tools.
He recognised the importance of rules and schemas, however.

Some of the Major AI centres

By the early 1970s there were several major AI centres, including at least

USA

MIT (Minsky & Papert...),
Stanford (McCarthy, Nilsson, ...,),
CMU (Newell, Simon, ...),
Yale (Shank, Abelson)

Edinburgh in the UK.
Hamburg and other places in Germany,
Rome in Italy.

and several other places in other countries and continents.

I was very lucky because Bernard Meltzer found money to get me to spend the year 1972-3 in the University of Edinburgh.

I was supposed to be doing an ambitious project to design a mind for a simulated robot called Adam, in a world called Eden.

But actually I was learning all sorts of things from a lot of very smart people –

I felt like a four year old child again.

When I returned to Sussex we started a new undergraduate programme, which included AI, Philosophy, Linguistics, Psychology. I think it was one of the first in the world.

A problem for current researchers

I had time to learn and experiment with new ideas and new techniques.

Nowadays that would be very difficult.

Unfortunately the pressure to publish and get grants now makes it very difficult for a young academic to spend so much time learning, after doing a PhD and getting a job: so people have to remain narrow.

This is partly a consequence of using performance metrics to evaluate individuals and determine funding allocations – as if doing research were like selling cars.

Politicians and university managers: take note!

What's the alternative?

1. Have very deep selection processes for university staff, and internal guidance and monitoring, using process-based evaluation, not performance-based evaluation.

2. Be more prepared to take risks with young researchers/teachers, especially if they are excellent teachers.

3. Perhaps base research funding on a weighted lottery scheme

   `http://www.cs.bham.ac.uk/~axs/lottery.html`

AI complemented old ways to study minds

Most AI researchers were doing engineering. But some wanted to do science, including understanding human minds –
e.g. Newell and Simon (Human Problem Solving 1972)

There are many old ways to study human minds:

- Reading plays, novels, poems. Many writers are shrewd observers! Reading their works will teach you much about how people see, act, have emotions, moods, attitudes, desires, etc. think and behave, and how others react to them.
- Studying ethology will teach you about how mental phenomena, including cognitive capabilities vary among different animals.
- Studying psychology will add much extra detail concerning what can be triggered or measured in laboratories, and what correlates with what.
- Neuroscience teaches us about physiological brain mechanisms that support and modulate mental states and processes, and are modulated by them.
- Studying therapy and counselling can teach you about ways in which things can go wrong and do harm, and some ways of helping people.
- Studying philosophy (with a good teacher) may help you discern muddle and confusion in attempts to say what minds are and how mental states and processes differ from one another and from physical states and processes.

All can be vastly improved by adopting the design-based approach: try to design and implement working models.
Think about architectures, mechanisms, representations, processes.

Why philosophy needs AI: two Examples

Free will

Philosophical discussions about free will are often based on simplistic assumptions about the kinds of mechanisms that might support deciding.

This leads to spurious oppositions between determinism and freedom.

By exploring a wide variety of information-processing architectures for control systems, whether produced by evolution or by engineers and philosopher-designers, we can show that there are more varied and complex cases than philosophers had previously considered, and we can explain why desirable forms of freedom and responsibility (e.g. doing what you want) depend on deterministic mechanisms rather than being incompatible with them.

Other kinds of freewill, the theological and the romantic notions, are incoherent:

`http://www.cs.bham.ac.uk/research/cogaff/misc/four-kinds-freewill.html`

Consciousness

By investigating architectures involving multiple concurrent sub-architectures, including some that monitor and modulate others, we can begin to understand more varieties of consciousness and self consciousness than philosophers are able to dream up in their arm chairs.

Is a fly conscious of your hand approaching when you try to swat it?
Is an operating system conscious of user attempts to violate file access restrictions?

The ordinary language concept ‘consciousness’ is not sufficiently precise to be used to formulate scientific questions!
AI researchers should model more specific things, e.g. attention, inference, wanting, noticing....

Why AI needs philosophy 1

AI needs conceptual analysis

AI researchers are often insensitive to the crudeness of the questions they ask

e.g. some ask 'How can we model emotions?' unaware that they are muddling up motivations, values, tastes, preferences, ideals, inclinations and many different sorts of things that do not necessarily involve being emotional. (Show Emotions demo if there's time.)

Analysing different mental concepts carefully shows that different mental phenomena presuppose different kinds of architectural complexity.

E.g. thinking about someone else's motives requires an architecture that includes the ability to represent mental states of others.

This requires meta-semantic competence: the ability to represent things that represent something else. (Includes handling referential opacity.)

That is also required for shame, e.g. being ashamed of your own motives.

Similar comments can be made about claims to model learning, creativity, consciousness: they all have complex presuppositions that lead to architectural requirements.

Another example is causation: our use of the word 'cause' is very subtle and complex and very hard to analyse. So if you say a system understands causation or learns about causation analysing that claim needs great care. `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506`

Researchers who assume an over-simple analysis, may end up making inflated claims (e.g. claiming to have modelled learning, or emotions, or scientific discovery or causal reasoning, when all they have modelled are very simple and shallow special cases).

Digression (20/06/2006): How to analyse concepts

Don't just take some word (e.g. 'emotion', 'consciousness', 'intention', 'attention') and try to find out what it means. Even the professionals disagree, and any definition you find is likely to be shallow and mistaken. Instead always work with families of related words.

E.g. don't just consider 'anger'. Put it alongside other examples of negative affect, e.g. 'irritation', 'annoyance', 'rage', 'outrage', 'disgust', 'regret', 'grief', 'disappointment', 'shame', 'guilt', 'embarrassment', 'dismay', 'fear', 'suspicion', 'grievance', 'wanting revenge'.

Then try to devise a collection of scenarios for which these words (in noun, adjective, adverbial, or verbal forms) would be appropriate or inappropriate. For each apparently closely related pair, try to find scenarios where one is appropriate and not the other.

Then try to work out what design features in the people (or animals) are required for those scenarios, and which design features are needed for one scenario and not another.

Example: (a) Your child is ill, and caring for her makes you miss an appointment. (b) Your child breaks something in order to gain your attention, and dealing with the breakage makes you miss an appointment. In each case are you angry? Irritated? Regretful? Disappointed? Wishing you had done something different earlier?

Harder: what architectural features support those states: what forms of representation are needed, what kinds of knowledge, what sorts of goals, preferences, values, intentions, what sorts of control mechanisms, what sorts of perceptual capabilities.

If you look at only one case (e.g. anger) or only one type of state (e.g. emotion) you are bound to ignore some of the important features and get things wrong.

It's like trying to produce a theory of evolution that explains only the evolution of butterflies.

Based on work by G.Ryle (e.g. The concept of mind) and J.L.Austin (e.g. 'A plea for excuses').
For a tutorial see `http://www.cs.bham.ac.uk/research/cogaff/crp/chap4.html` (1978)

Why AI needs philosophy 2

Many of the best AI researchers in the first decade expected that all reasoning or problem solving could make use of essentially logical or sentential information structures, and logicist AI has many important achievements.

I had been doing a philosophical analysis of what reasoning is, e.g. in mathematics.

namely manipulating an information structure in such a way as to preserve some aspect, e.g. truth or denotation

All birds are mortal

All chickens are birds

Therefore: All chickens are mortal

Manipulating diagrams can also preserve denotation or truth.

I used this as an attack on logicist AI in my first AI paper, in 1971 (2nd IJCAI).

I argued that both Fregean and and analogical modes of representation and reasoning were important and could be useful, in different sorts of problems.

This is also relevant to the nature of mathematics.

Many others have made the same point, but there has been little success in modelling visual/spatial/diagrammatic reasoning: mainly because most of the problems of vision are still unsolved in AI, even though there has been a lot of work on sub-problems, such as recognition, tracking and route-finding.

There is far more to seeing a spanner than recognising it, as you can tell by watching a 3-year old trying to use one.

Other philosophers

Two important philosophers whose interest in AI grew in that period were

Dan Dennett, whose book Brainstorms (1978) also attempted to build bridges between the two disciplines

Margaret Boden whose two books (Purposive Explanation in Psychology (1972) and Artificial Intelligence and Natural Man (1978)) helped to spread the word to wider audiences.

[OUP will shortly publish her new 2 volume History of Cognitive Science which will help to illuminate the early years of AI.]

Other philosophers also became interested, but not many – not nearly as many as I thought would. My prediction in 1978:

within a few years philosophers, psychologists, educationalists, psychiatrists, and others will be professionally incompetent if they are not well-informed about these developments.

Many philosophers still remain pretty ignorant about computing and AI at the end of their PhD studies, even if they have managed to learn to use Word, Powerpoint and web browsers.

In that book I also predicted that learning to design, debug, document complex working theories, would transform education, and general self understanding.

Alas the one thing people, including most school kids, don't get on a typical PC is any software development tool, e.g. compiler, interpreter for nice high level programming language: a dreadfully wasted educational opportunity. (Compare the BBC micro.)

Steps towards architectures

During those early years it became clear that whereas much of AI research in the past had been focused on algorithms and representations, it was also necessary to start thinking about how to put all the pieces together in an architecture combining multiple kinds of functionality, in concurrently active components, especially if we are to explain or model the kind of autonomy and creativity found in humans and other animals.

This was specially obvious to anyone who was trying to use AI to understand aspects of human minds, since it had been clear that human minds are multi-faceted systems with many components concurrently active.

The idea that we use a sense-decide-act cycle should have been obviously false to everyone, but wasn't for some reason.

The need for an architecture, and specification of some requirements, was the topic of Chapter 6 of The Computer Revolution in Philosophy (1978) now online

`http://www.cs.bham.ac.uk/research/cogaff/crp/`

However work on architectures integrating diverse components of a robot did not develop seriously until at least 10 years later.

Unfortunately, around that time people started convincing themselves that an insect-like architecture would suffice for intelligent systems, and research was held back for years – with many young researchers given false beliefs and vain hopes.

Most implemented robot architectures are still very primitive, compared with a human.
Or even an insect.

A generative framework for describing architectures

We don’t need one architecture. We need to understand options and tradeoffs.

We need to be able to classify types of components in a principled way and talk about varieties of relationships between components.

The CogAff schema is a first draft simple example:

9 main types of concurrently active components, with many possible links between components.

This can accommodate a wide variety of types of architectures.

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

It would help to have a widely used ‘grammar’ for types of architectures instead of everyone inventing their own labels and diagramming conventions, making comparisons very difficult.

NB: an architecture need not be fixed: a human infant has an architecture that grows into an adult architecture that has many extra components.
Maybe some components used for bootstrapping are later discarded.

An example: Omega Architectures

An 'Omega' architecture has
– 'peephole' perception and
– 'peephole' action,

as opposed to

– 'multi-window' perception and
– 'multi-window' action.

In multi-window perception the perception systems include layers of abstraction that communicate directly with 'higher level' central systems.

Likewise multi-window action: gesturing and speaking are actions that are different in many ways from posture adjustments.

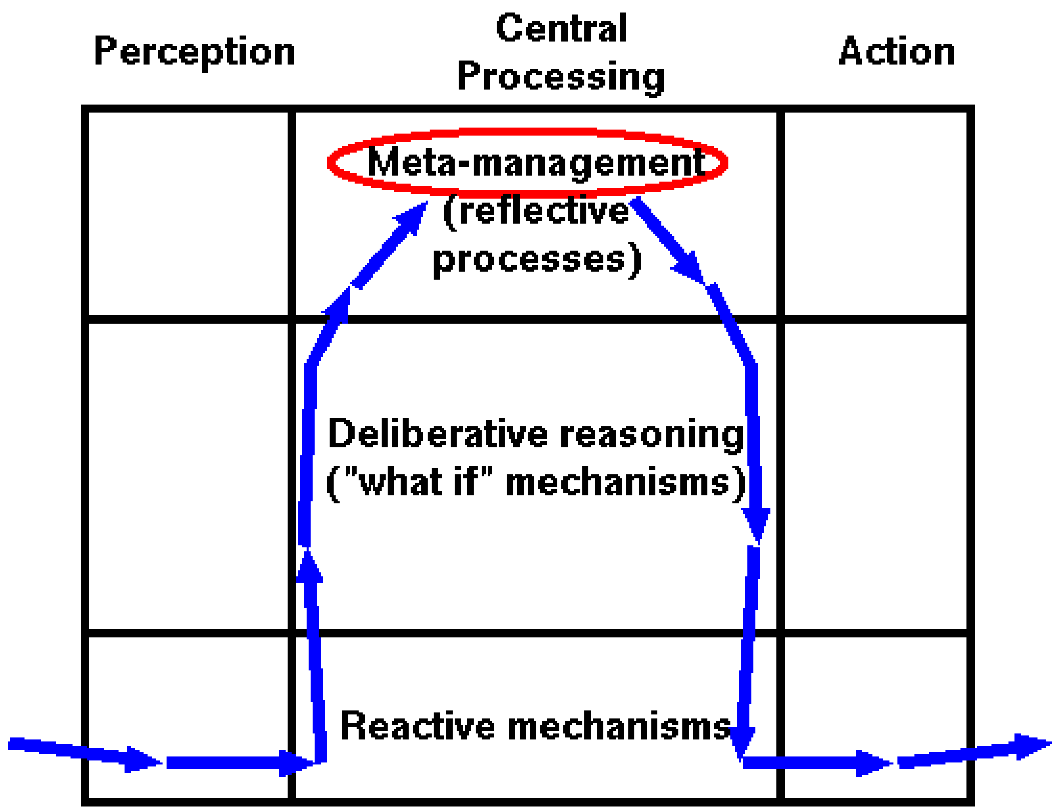

As far as I know very few AI systems, if any, have multi-window perception and action: it's mostly all peephole stuff.

Another example: Subsumption

In most subsumption architectures the top two CogAff layers are nonexistent, and all layering is done within the reactive level.

Adding deliberative and meta-management layers enormously enriches and transforms subsumption.

Why?

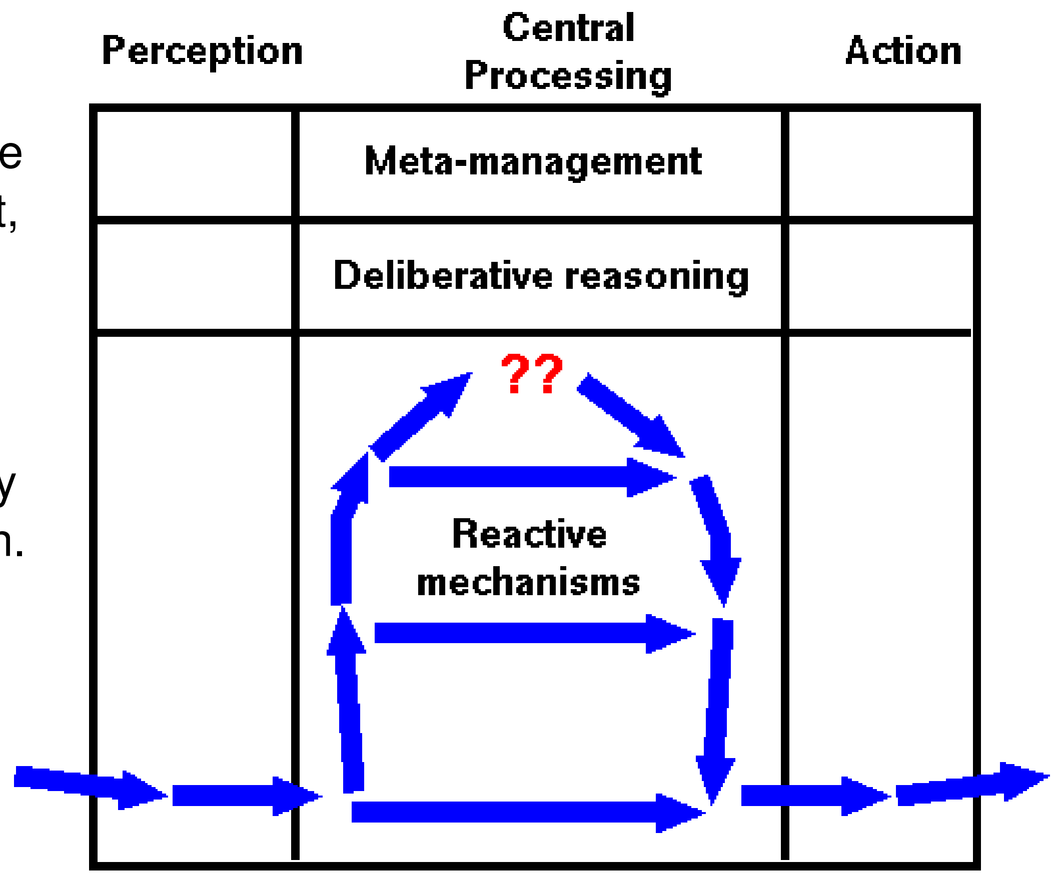

An architecture based on conceptual analysis

H-Cogaff: a conjectured adult-human architecture (bird's eye view)

A special instance of the CogAff schema is the H-CogAff architecture, crudely depicted on the right – and conjectured to represent some important aspects of a normal adult human (with much detail missing).

In practice there are likely to be far more connections between components of the architecture than shown by the arrows.

The architecture has to grow itself:

An infant does not start off like this! See Minsky's new book: The Emotion Machine (available online)

and papers and presentations on the Birmingham Cogaff web site:

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

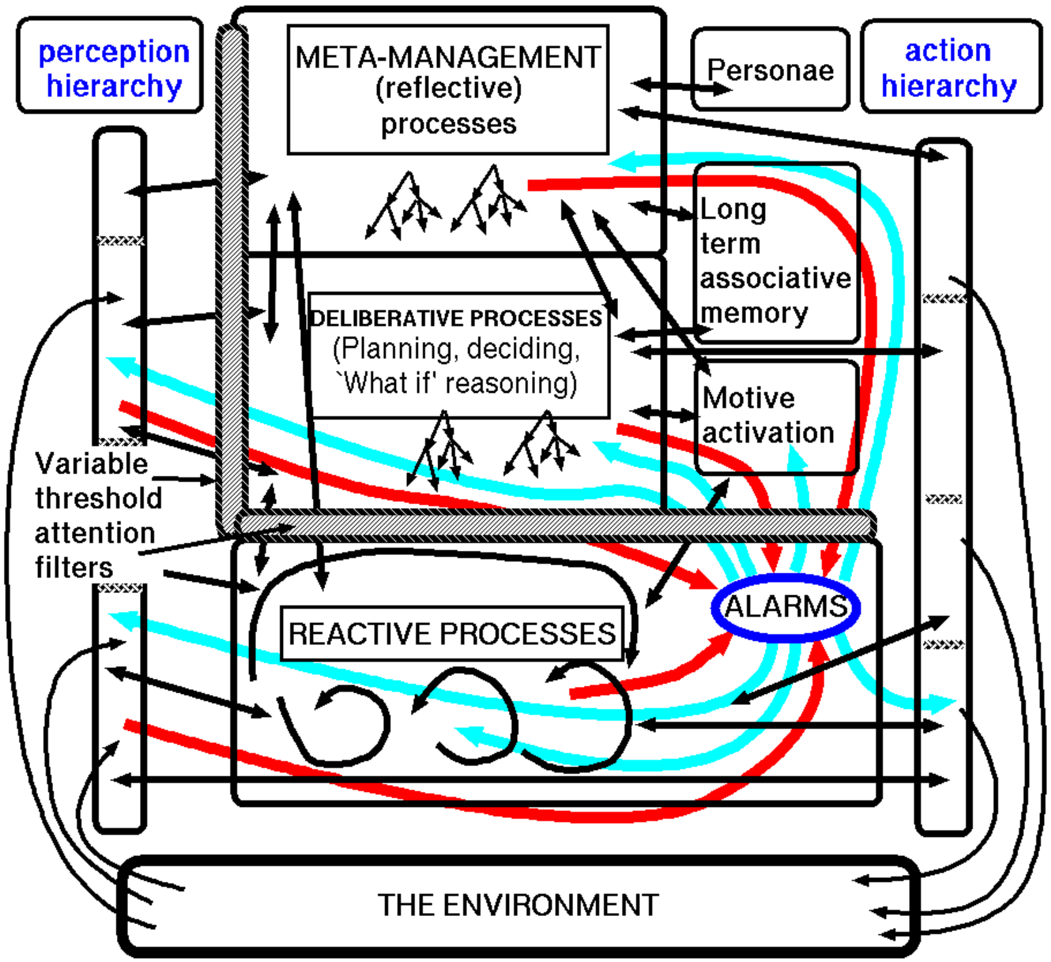

Design space and niche space

A biological niche is not a geographical location: it is more like a set of requirements for a design:

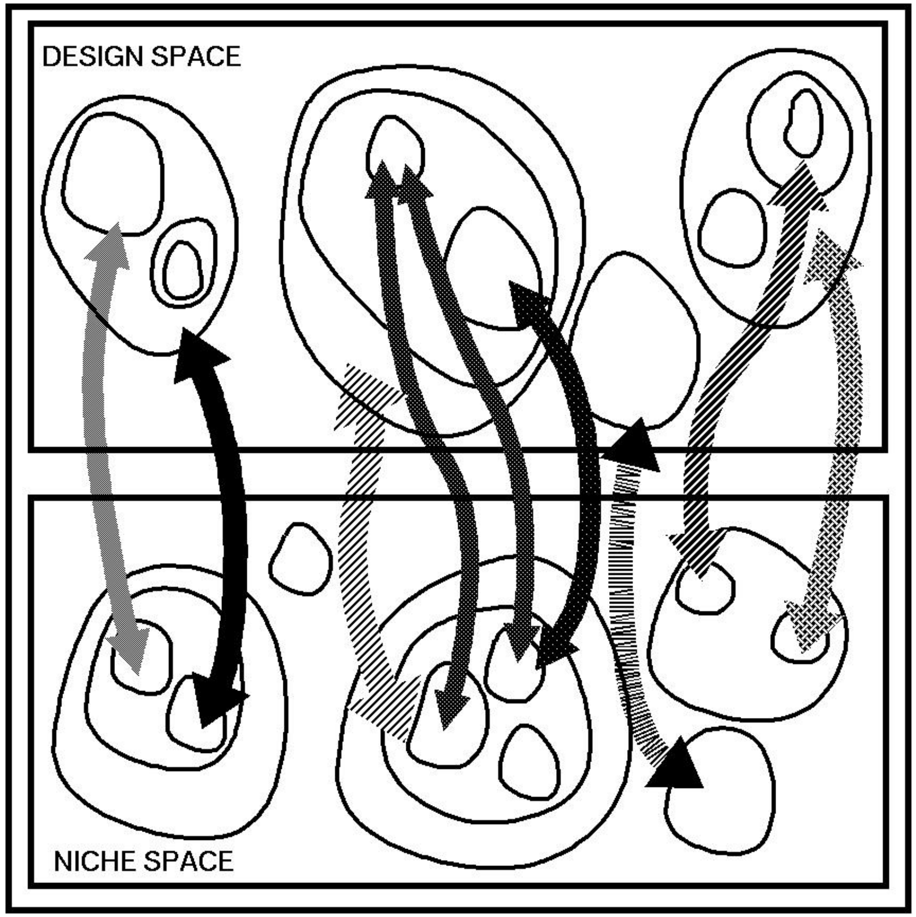

The arrows represent different kinds of fitness relations between types of designs and types of niches: analysis still to be done

We need to understand trajectories for evolution and for development: very complex dynamics.

Different sorts of trajectories when designs and niches change:

- individual development and learning
- evolutionary development of a species
- cultural/social development of species or groups
- 'repair' by an external designer/engineer

  The last may include big gaps in trajectories, while the others allow only small discontinues in design.

  Not all designs can cope with big discontinuities in niche-trajectories.
  (Humans can, better than most.)

Understanding the feedback loops in these trajectories may require new mathematics.

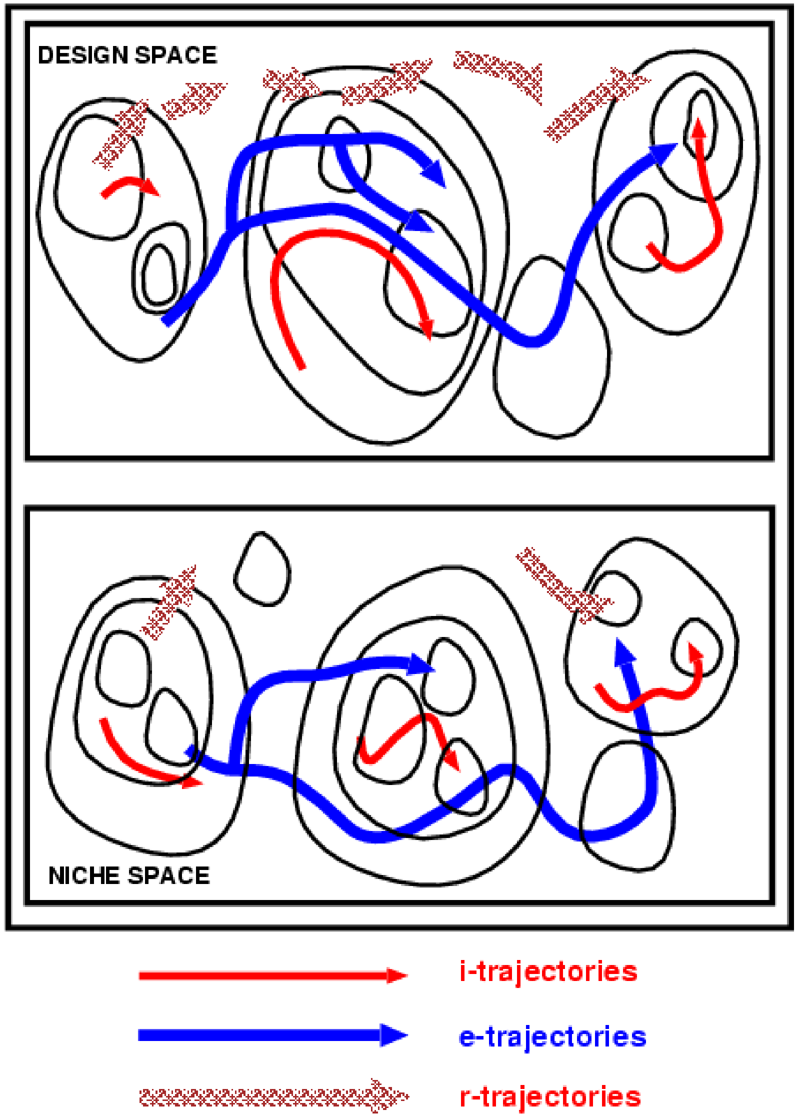

Philosophy inspiring AI

One of the reasons for AI researchers to learn philosophy is that old philosophical problems can inspire new AI research.

One example is the old philosophical debate between empiricists (e.g. Hume) and apriorists (mainly Kant).

We can now reformulate the debate in terms of investigations of nature-nurture tradeoffs.

Unfortunately many AI theorists just assume that any learning system must start off with as little prior knowledge as possible, and must derive all its concepts by abstraction from experienced instances (concept empiricism/symbol grounding)

as if proposing that the human genome should discard millions of years of learning about the nature of the environment: unlike all the many animals that start off highly competent at birth (precocial species: e.g. deer run with the herd soon after birth).

The time is ripe for AI researchers to re-open that discussion in collaboration with biologists studying varieties of animal cognition.

See paper in IJCAI 2005 on the ‘altricial precocial spectrum’ for robots.

It will also help to transform philosophy and developmental psychology.

(Show Betty, the hook-making Crow)

Kant vs Hume on Mathematics

A philosophical conflict between two philosophers, David Hume and Immanuel Kant drove my own research interests concerning the nature of mathematics.

Hume claimed that there is no kind of knowledge apart from the empirical knowledge acquired through the senses and trivial tautologies that are true by definition (sometimes called 'analytic' truths).

Everything else, he claimed, was nothing but 'sophistry and illusion' and should be consigned to the flames (e.g. theology and metaphysical philosophy).

Kant thought mathematical discoveries were not empirical and truly expand our knowledge.

Even things like $7 + 5 = 12$

He was right of course. To understand why, we need to model what goes on when a child learns about numbers.

AI work on different ways in which a machine could learn mathematics and derive new conclusions from old, including the use of analogical representations in some case, can help to show that Kant was right.

But such research is still in its infancy.

E.g. it is very difficult to give a computer an intuitive understanding of continuity.

But we need to do a conceptual analysis of that notion as part of the process.

Towards modelling a child learning mathematics

Most work on mathematical reasoning in AI has attempted to give machines the ability to do things adult mathematicians do.

While I was learning about AI, I watched our four year old son learn about numbers.

I came to the conclusion that there is an opportunity to learn something new and deep about human minds by trying to understand the early stages of learning mathematics by designing a child-like learner.

In Chapter 8 of CRP I tried to summarise some of the required capabilities

- Learning to sort things into groups
- Rhythmically performing a sequence of actions, e.g. pointing at objects, climbing up stairs, moving objects from one container to another.
- Learning to generate number names rhythmically.
- Learning to do both tasks in synchrony (including detecting and correcting lapses).
- Learning different stopping conditions corresponding to different tasks, e.g.
  - Run out of things to count,
  - A target number hit
  - A target arrangement hit
- Learning to apply counting operations to counting operations

  See: `http://www.cs.bham.ac.uk/research/cogaff/crp/chap8.html`

Studying such things may lead both to understanding better what needs to go on in a robot with human-like intelligence, including mathematical intelligence, and also understanding better what goes wrong in much mathematical education in primary schools because it is based on incorrect models of learning and discovery.

More on the 1970s

There were many technical achievements in AI in the 1970s, many of them concerned with new engineering applications including the early development of expert systems and many tools now taken for granted by researchers (e.g. Matlab, Mathematica).

A major robotic achievement, now generally forgotten, was Freddy the Edinburgh robot, which could assemble a toy wooden car in 1973, though it could not see and act at the same time, because of low computing power. Minsky's frame-systems paper was very influential, and inspired many formalisms and toolkits (also aspect graphs?). Logic programming started to take off.

AI vision research was also starting to get off the ground, at last moving away from pattern recognition. E.g. pioneering work was done by Barrow and Tennenbaum, published in 1978, and by others working on ways of getting 3-D structure from static or moving image data.

However many did not appreciate the importance of the third dimension and merely tried to classify picture regions – a task that still occupies far too many researchers who could be doing something deeper.

Gibson's ideas were just beginning to be noticed around that time, especially his emphasis on the importance of optical flow and texture gradients, and later his ideas about affordances.

Some people were already trying to resurrect neural nets, with limited success – and much optimism.

Many worked on new higher level languages and toolkits (though not architecture toolkits?).
Prolog took off – especially in Europe. There was much work on natural language processing, including European translation projects and the DARPA speech understanding project.

My own vision project (POPEYE) based on a multi-level multi-processing visual architecture made some progress then hit a funding wall. I also started trying, without much success, to get people to think about surveying spaces of possibilities and the tradeoffs therein instead of (vainly) competing to find the single best solution to a problem.

Freddy the 1973 Edinburgh Robot

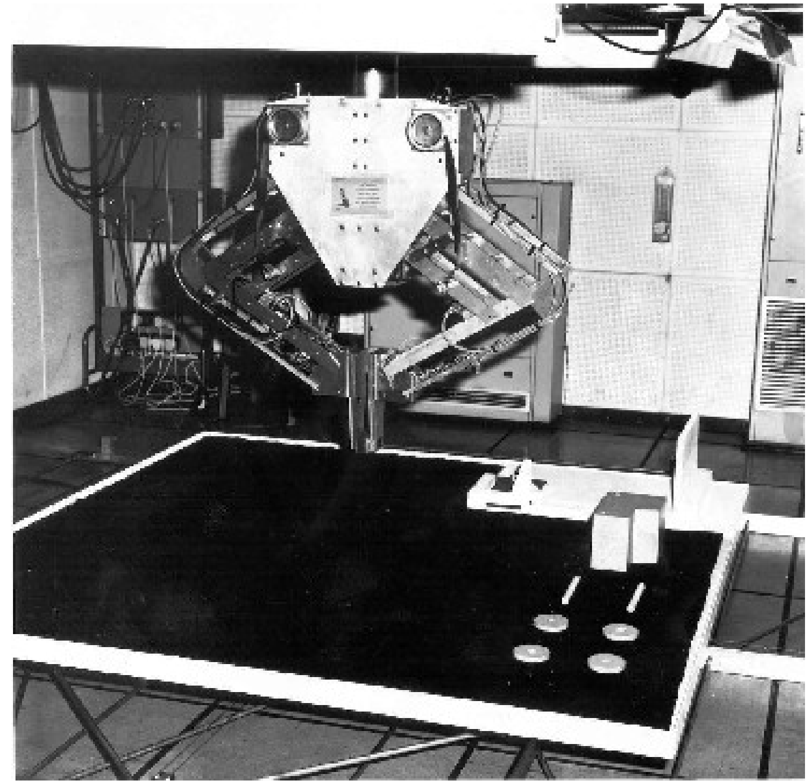

Freddy, the 'Scottish' Robot, was built in Edinburgh around 1972-3.

Freddy II could assemble a toy car from the components (body, two axles, two wheels) shown. They did not need to be laid out neatly as in the picture.

However, Freddy had many limitations arising out of the technology of the time.

E.g. Freddy could not simultaneously see and act: partly because visual processing was extremely slow.

Imagine using a computer with 128Kbytes RAM for a robot now.

There is more information on Freddy here

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

`http://www-robotics.cs.umass.edu/ARCHIVE/Popplestone/home.html`

In order to understand the limitations of robots built so far, we need to understand much better exactly what animals do: we have to look at animals as engineers, asking, repeatedly:

How could we design something that works like that?

`http://www.cs.bham.ac.uk/research/cogaff/misc/design-based-approach.html`

On seeing at multiple levels

The POPEYE program (1975-1978), summarised briefly in chapter 9 of The Computer Revolution in Philosophy, processed different levels of interpretation of a complex noisy picture in parallel, using concurrent top-down and bottom-up processing and different ontologies.

Not to be confused with 'heterarchy' where a single locus of control moves between different subsystems: that is much less robust and flexible.

Similar things were being done with speech processing – e.g. in the DARPA speech understanding project.

I still think visual processing involves multiple levels of processing going on concurrently, but the contents of perception are primarily processes at different levels of abstraction, since normally perception is of a changing environment, not recognition of a picture, or keyhole view of a static scene.

Structures are then perceived in the context of processes involving them. Many structures are flexible or articulated, so that they too are processes.

A structure is a special case of a process where all velocities are 0.

Further developed in:
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505`

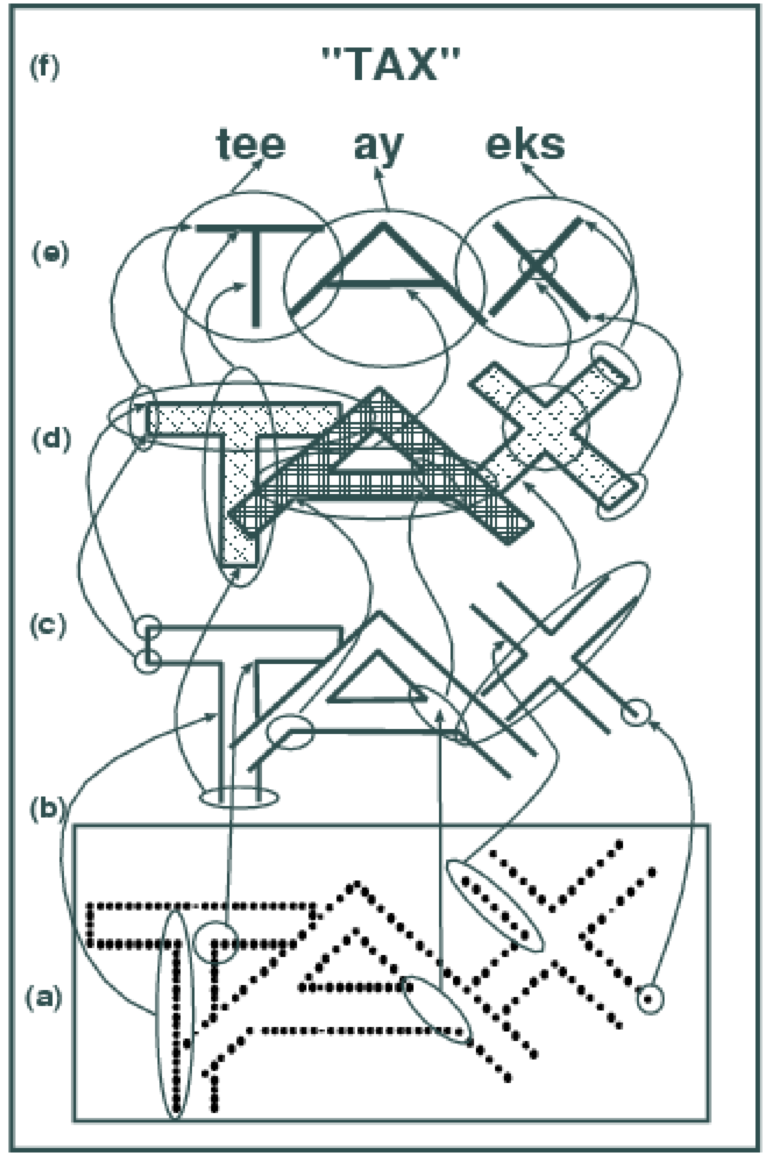

The need to reassemble AI

During the following decade the field started increasingly to fragment for several different reasons (including rapid growth in numbers), with many bad effects, including killing off some major promising developments (e.g. research on 3-D vision required for manipulation).

AI has become far more a collection of narrow specialisms with most researchers barely aware of anything going on outside their own sub-fields.

there has been much fragmentation, within each of: AI, psychology, neuroscience — most researchers focus only on a limited sub-field, e.g.

- vision (usually low-level vision nowadays)
- language (text, speech, sign-language)
- learning (many different kinds)
- problem solving
- planning
- mathematical reasoning
- motor control
- emotions

  etc....

It is not obvious that systems developed in that way can be combined with other parts of an integrated working robot.

(See scaling up vs scaling out, below.)

Towards a new integration

Perhaps we can now start re-integrating AI, both as engineering and as the most general science of mind.

At least the hardware support is more powerful than ever before.

For a ‘Grand Challenge’ proposal see

`http://www.cs.bham.ac.uk/research/cogaff/gc/`

And for a suggested means to re-integrate AI see

`http://www.cs.bham.ac.uk/research/cogaff/gc/aisb06/sloman-gc5.pdf`

Building roadmaps for AI research:
Analyse complex tasks and work backwards through a partially ordered network of simpler scenarios, till you get to something you could start working on.

Beware tempting dead-ends that will not lead where you want to go (even if you can demonstrate improvement on some benchmark).

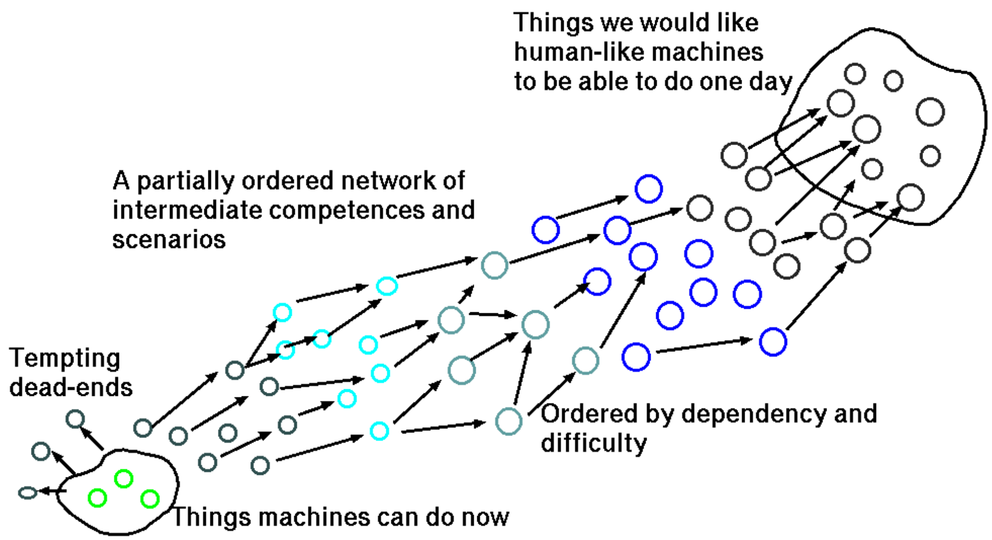

Note the overlap between engineering (achieving complex practical goals) and science (explaining complex natural phenomena).

Current studies of mind

- Current ways of studying (animal, human and robot) minds are
  - too fragmented
  - too riddled by turf wars
  - too much influenced by prejudice (what people would like to be true)
  - based on inadequate notions of science and explanation
  - based on too little data in forms that are too restricted, or too much data of the wrong sort

- Examples:
  - bad theories about emotions
  - confused concepts treated as well understood
  - theories/models/explanations that don't 'scale out' (fit into a larger context)

- We can remedy this by working out the implications of these facts:
  - minds DO things: they are constantly active machines
  - there is not just one kind of mind: very many exist in nature, even among humans: young, old, normal, damaged, ancient, modern (industrialised)
  - all organisms are information processors
  - evolution is far ahead of our understanding
  - all complex designs involve complex trade-offs
  - new evolutionary designs do not simply throw away old solutions, but build on them:
    humans share much with much older species

Scaling out vs scaling up

The need to 'scale out' (combining with other capabilities)

is at least as important as

the need to 'scale up' (coping with complexity in a problem)

There is no guarantee that a technique, or form of representation, or algorithm, etc. that works for an isolated task will also work when that task has to be integrated with many other kinds of functionality in an integrated system.

This is true of AI techniques that 'scale up' very well within a particular application domain, e.g. path planning.

E.g. they may not 'scale out' to support anytime planning or reasoning about planning, or cooperative planning, or explaining a plan while developing it, or coping with new visual information relevant to an incomplete plan.

Human abilities generally do not scale up: we are defeated by combinatorics.

(Donald Michie referred to 'the human window'.)

But they scale out and interact fruitfully: e.g. what you see can help you understand words you hear and vice versa. (McGurk effect)

Babies, blankets and string

Scaling out could be demonstrated in scenarios like this.

Learning how to get hold of a toy that is out of reach.

Contrast:

- Short blanket
  - Grab edge and pull
- Long blanket
  - Repeatedly scrunch and pull
- Towel
  - Like blanket
- Sheet of plywood
  - Pull if short, otherwise crawl over or round
- Sheet of paper
  - Roll up? (But not thin tissue paper!)
- Slab of concrete
  - Crawl over or round
- Taut string
  - Pull
- String with slack
  - Pull repeatedly
- String round chair-leg
  - Depends
- Elastic string
  - ?????

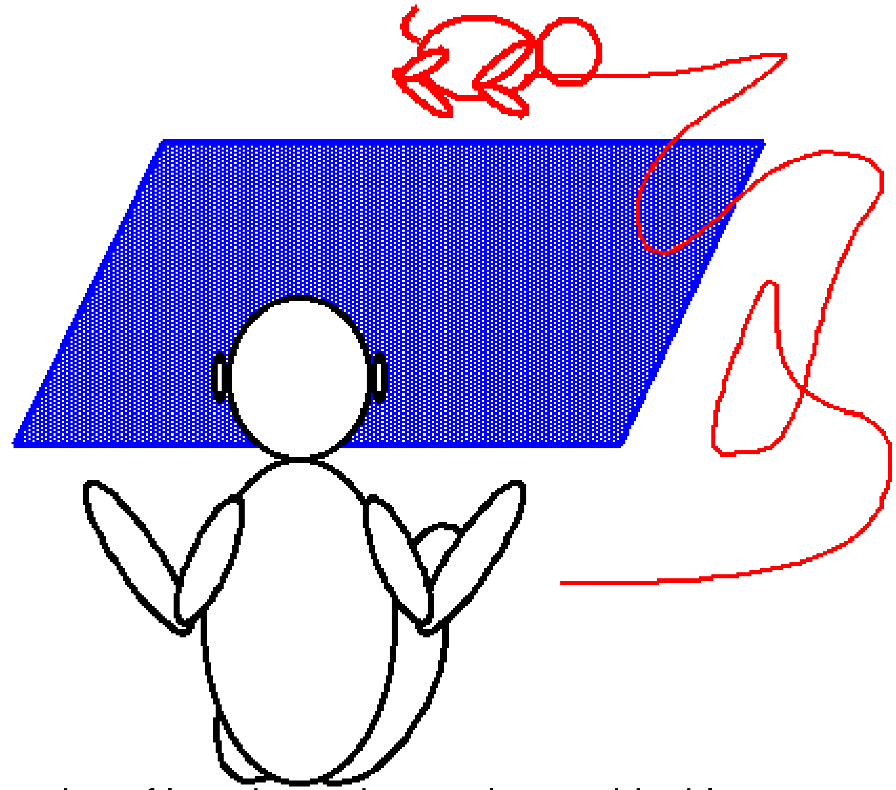

See this discussion of learning orthogonal recombinable competences
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601

On seeing manipulable things

Despite the poor quality of the image, you can probably see many points at which you can touch or grasp the objects in this scene. You can also work out (roughly) at which angles you would need to orient your fingers and the direction of approach required in order to achieve the different tasks.
You can probably also visualised, at least crudely, some actions you could perform to bring about a situation where the spoon is on the saucer and the saucer is on the cup, upside down.

Can your current robot do that?

Probably not. Even if it recognises cup, saucer and spoon: much easier than seeing surface structures and affordances.

See
`http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf`

More visual challenges:
`http://www.cs.bham.ac.uk/research/cogaff/misc/multipic-challenge.pdf`

Using factual material

- One problem is identifying what needs explaining.

  Too often people observe only what their theories deem relevant, or collect only information that their statistical tools can process.

- A scenario-based approach can help to overcome that limitation

  by collecting and analysing very many real scenarios, organised according to their similarities and differences and ordered by complexity
  e.g. (of mechanisms, of information, of architectures, of representations needed).

Examples: collect and study videos of animals and children:

- Betty, the new caledonian crow, surprised researchers at the Oxford University Zoology department when she displayed an ability to make a hook out of a straight piece of wire, in order to fish a bucket containing food out of a tube:

  (`http://news.bbc.co.uk/1/hi/sci/tech/2178920.stm`)

- An 18 month old child attempts to join two parts of a toy train by bringing two rings together instead of a ring and a hook, and showing frustration and puzzlement at his failure. (`http://www.cs.bham.ac.uk/~axs/fig/josh34_0096.mpg`)

  A few weeks later he was able to solve the problem: what had changed?

- If time: video of the child playing with trains on the floor about a year later.

Supplement observed scenarios with a large collection of analytical scenarios: compare Piaget

Two-way scientific information flow

We need a far better understanding of how natural intelligence works, at different levels of abstraction, if we are to build more intelligent (e.g. robust, autonomous, adaptive) artificial information-processing systems.

In particular, building working human-like robots requires us to develop architectures combining many types of functionality.

But in order to understand examples of natural intelligence we need to understand how to design systems with similar capabilities.

Emotions and control mechanisms

What is there in common between

- a crawling woodlouse that rapidly curls up if suddenly tapped with a pencil,
- a fly on the table that rapidly flies off when a swatter approaches,
- a fox squealing and struggling to escape from the trap that has clamped its leg,
- a child suddenly terrified by a large object rushing towards it,
- a person who is startled by a moving shadow when walking in a dark passageway,
- a rejected lover unable to put the humiliation out of mind
- a mathematician upset on realising that a proof of a hard theorem is fallacious,
- a grieving parent, suddenly remembering the lost child while in the middle of some important task?

Proposed Answer (not original – e.g. see Herb Simon on emotions):

in all cases there are at least two sub-systems at work in the organism, and one or more specialised sub-systems can somehow interrupt or suppress or change the behaviour of others, producing some alteration in (relatively) global (internal or external) behaviour of the system — which could be in a virtual machine.

Some people would wish to emphasise a role for *evaluation*: the interruption is based at least in part on an assessment of the situation as good or bad.

Is a fly capable of evaluation? Can it have emotions? Evaluations are another bag of worms.

Some 'emotional' states are useful, others not: they are not required for all kinds of intelligence — only in a subset of cases where the system is too slow or too uninformed to decide intelligently what to do — they can often be disastrous!

See: `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cafe04`

Do machines, natural or artificial, really need emotions?

Lots more to be said but....no time

Join the 300 year project.

(Don't believe claims about what's imminent.)

Very difficult, but enormous fun.

Learn more about what you are.

Many potential applications too – not just smart machines but smarter ways of dealing with people

(e.g. in school, in therapy, counselling).

Help to revolutionise several old disciplines.

Maybe even computer science?

Comments, criticisms and suggestions always welcome.

`http://www.cs.bham.ac.uk/~axs/`

I may extend these slides later, to fill some of the gaps.

See also Minsky's web page: The Emotion Machine is online.

`http://www.media.mit.edu/~minsky/`

Lots more here: `http://www.cs.bham.ac.uk/research/cogaff/talks/`