Cafe Scientifique Launch
Midlands Arts Centre
25 October 2002

WHAT IS SCIENCE?

(Can There Be a Science of Mind?)

Aaron Sloman

School of Computer Science,
The University of Birmingham, UK

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

These slides are available as talk 18 here

http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk18

also on Slideshare.net http://www.slideshare.net/asloman/presentations

The context: Cafe Scientifique

For some years there has been an informal “Cafe Scientifique” in Leeds at which people interested in learning and talking about science meet for a meal or drinks, hear a presentation by an invited speaker and join in discussion and debate.

There is now a project to promote similar things in many parts of the UK with a national coordinator sponsored by the Wellcome Trust.
Details can be found at this web site:

`http://www.cafescientifique.org/`

Dr Emil Toescu in the Medical School, University of Birmingham, arranged a “launch” meeting for a Birmingham Cafe Scientifique on Friday 25th Oct 2002 at the Midlands Arts Centre.

`http://www.macarts.co.uk/`

I was asked to introduce the discussion and debate and suggested ‘What is science?’ as a topic. Hence these slides.

For more information about Cafe Scientifique meetings in Birmingham see

`http://www.cafescientifique.org/birmingham.htm`

ACKNOWLEDGEMENT

This talk was prepared using only
reliable, portable, free software,
e.g. Linux, Latex, ps2pdf, gv, 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

WHAT ISSCIENCE?

?????
?????

WHAT IS SCIENCE?

A common answer:

Science is
A search for the laws of nature???

We can improve on that answer!

Newton's achievement

Newton was one of the greatest scientists of all time, although, as he said, he "stood on the shoulders of giants".

E.g. Galileo helped provide the basis of Newtonian mechanics, and Descartes' arithmetization of geometry was an essential precursor to Newton's achievements.

Among many other things, Newton invented three simple laws whose explanatory and predictive power turned out to be astonishing:

1. Every object in a state of uniform motion tends to remain in that state of motion unless an external force is applied to it.

   In other words, every physical object has inertia.

2. The relationship between an object's mass m, its acceleration a, and the applied force F is $F = ma$.

   Acceleration and force are construed as *vectors*, i.e. they have direction and magnitude, and the direction of the acceleration is the same as the direction of the force.

   Note that this implies the first law. If the mass is non-zero and the force is zero, then the acceleration must be zero.

3. For every action there is an equal and opposite reaction.

There are many excellent internet web sites explaining this and other aspects of the history of science in detail, e.g.

`http://csep10.phys.utk.edu/astr161/lect/history/newton.html`

But the laws are icing on the cake

Notice that in order to be able to formulate those three laws, Newton had to use concepts that refer to types of entity that can exist in our universe, e.g.

- Physical objects
- The velocity of an object
  This assumes that we can also refer to
  - The position of an object
  - Different times (at which objects can have different positions and other properties)
- The acceleration of an object
  which presupposes that velocity is something that can change
- The force acting on an object
  which itself may be the resultant of many forces
- The mass of an object
  often confused with weight, which is just the force with which the object is pulled towards the centre of the earth (a force that varies with the location of the object).
  Mass is something more subtle – resistance to change of velocity.

So in order to be able to formulate his laws Newton presupposed an ontology of objects, properties, relationships, processes, causes.

(Actually the presuppositions were more precise than that, e.g. it was presupposed that many of the properties and relations could be mapped onto the continuum of real numbers, a notion that was not analysed fully until the 19th century!)

What is an ontology?

Roughly: it is a collection of concepts referring to things that can in principle exist, plus some sorts of composition rules specifying how more complex entities can be formed from simpler ones.

- Exactly what this means is a complex and subtle question.
- Different animals use different ontologies.
- Young children have different ontologies from older children, whose ontologies are different from adult ontologies.
- People in different cultures have different ontologies.
- A culture's ontology can change over time: in Aristotle's Greece nobody had an ontology including electromagnetic radiation.
- Most special sciences use ontologies that build on and extend the ontologies used generally in the culture.

  E.g. Biology adds genes, genome, chromosomes, species, selection, niches, epigenesis ....
- Often a special notation or branch of mathematics has to be developed to express the concepts, e.g. notations for chemical compounds, or for grammars in linguistics.
- A result of the development of computers, computer science, artificial intelligence and software engineering is that some people now use an ontology containing virtual machines, compilers, data-structures, interrupt-handlers, network protocols, software architectures, viruses ....

  See: `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#wpe08`

What you can do with an ontology

Having got his ontology Newton could ask questions that could not be asked without it.

- Can an object have property X without having property Y?

  E.g. can an object have zero force on it without having zero acceleration?
  No.

- Can property X vary without property Y varying?
  E.g.
  - Can the position of an object change without a force acting on it?
    yes – if it is already moving, it will continue moving.
  - Can the velocity of an object change without a force acting on it?
    No
  - Can the acceleration of an object change without the force acting on it changing?
    No – unless its mass changes (impossible for Newton).

- If one change causes another how are the changes related?

  Newton's second law $F = ma$ summarises this, but in order to deal with its effects over time, especially when F changes, he had to invent a new kind of mathematics, i.e. differential and integral calculus (also invented by Leibniz), to work out how to deal with forces (and other things) that change over time.

  So Newton's laws are answers to questions about what is and is not possible – questions that could not even be formulated without the appropriate ontology.

QUESTIONS versus ANSWERS

Perhaps the most important thing achieved by developing a new ontology or extending an old one is:

being able to ask questions that could not previously be expressed.

For example, if you have an ontology which allows atoms to have nuclei composed of various combinations of protons and neutrons you can ask questions about which combinations are possible, e.g.

Is it possible for an atom to exist with a number of protons that has not yet been found in any atom?

Such questions can lead both to empirical investigations to find the answer and to theoretical investigations to see whether a good theory allows such a combination or rules it out.

Only the discovery of laws restricting possible combinations of sub-nuclear particles could rule it out.

We are now in a better position to answer the question: What is science?

WHAT IS SCIENCE?

A better answer – science is a disciplined attempt to find out:

- what exists

  E.g. people, fleas, clouds, rivers, atoms, sub-atomic particles, molecules, poverty, wars, minds, emotions, computational processes in computers, genes, species, niches, ecosystems...

- how things work

  E.g. how molecules of atoms and of hydrogen can be transformed into molecules of water, how centripetal force produces circular or elliptical motion, how an egg develops into a chicken, how humans generate grammatical sentences (sometimes).

- why they work as they do

  This usually requires appeal do a deeper theory. Usually mathematics is required to derive precise consequences from a deep theory.

- what doesn't exist but could exist

  Many animals that might have evolved did not. Many molecules might be produced that have not been. Could peace exist on earth?

- how such things would work if they existed

  What sort of mechanisms could make cold fusion happen? What kind of design would allow a computer to learn to talk and understand English fluently?

...continued

WHAT IS SCIENCE, continued...

.... a disciplined attempt to find out:

- what sorts of things cannot exist

  An element of atomic number 2000 ? A mouse proving theorems about algebra? Telepathy?

  Note: a law stating that whenever X and Y happen Z also happens is equivalent to a statement that it is impossible for X and Y to happen without Z also happening.
  I.e. laws state that something is impossible (or in some cases improbable).
  But they always presuppose a deeper theory implying that the things referred to as impossible together can exist separately, e.g. X and Y without Z.

- under what conditions they cannot exist

  Some laws are universal, while others (e.g. unsupported objects accelerate at 980 cm per second per second) work only in restricted contexts. Some are probabilistic, e.g. throwing two dice will produce two sixes only about 2.7778% of the time. What about: “a chess machine will not move its king into check”?

- why they cannot exist

  The only way to explain why something is impossible (why a law holds) is with reference to a deeper, more general theory. Einstein’s theory of general relativity provided an explanation for Newton’s law of gravitation.

But these answers are too general.

There’s much more to be said in detail about the nature of science.

Craft, Science, Engineering

Historically there is often a progression through craft, then science to engineering, which can produce new craft, new science and new engineering, indefinitely.

Craft:

We develop skills, learn from experience, teach others, solve many practical problems.

But we don't really understand why our methods sometimes work and sometimes fail.

Science:

We describe systematically and explicitly what was previously only intuitive.

This often requires mathematics (not necessarily numbers – e.g. grammars or program specifications).

We link different theories together to form systems, some of which explain others.

We accept nothing on trust, have no faith about anything, but we persist when it looks as if there's more to find out. A scientist should never believe anything as **proven**, but can **provisionally decide** that one theory is currently better than all rival theories.

We publish theories and data, and we invite and attend to criticism (or should do!)

Engineering:

We use the science to refine, explain and extend what was previously only craft,

e.g. explaining why old bridges are stable, designing a new kind of bicycle frame.

We use engineering advances to probe nature in more depth and with greater precision, discovering more for science to explain, and sometimes falsifying theories in new ways.

Components of theories

For every kind of scientific subject matter we need

- A form of representation (often using some kind of mathematics)
- An ontology (catalogue of kinds of things that can exist)
- Techniques and tools for manipulating the representations, so as to model things, draw inferences, and sometimes replicate.

A lot more needs to be said about the forms of representation used in various sciences.

- Often the development of a new notation or form of representation, together with a new body of mathematics for reasoning with it, is crucial to an advance in science.
- For example, the development of new programming languages helps to advance the science of information processing.
- New forms of representation that are good for science sometimes also advance engineering, and vice versa.

Blurred boundaries

The boundaries between craft, science and engineering are not sharp.

- The craftsman or artisan who reflects on what does and does not work, keeps records, and looks for patterns, is already moving towards being a scientist.
- And sometimes creative scientists are not terribly disciplined; they have hunches, they have prejudices against certain theories, they may fail to grasp a new unfamiliar ontology.
- Sometimes a high-priest mentality, or intellectual snobbery leads people to think that only what they do can be called science: they may be unaware that their rigid constraints (e.g. use only numerical data that can be fed into statistical packages to produce significant correlations) may be obstructing deep science.
- In particular, some non-physicists mistakenly believe that the essence of physics is collection of measurements and a search for laws consistent with the measurements; so they teach their students that that is the *only* way to do science: corrupting the minds of the young in the name of science.

We'll see later that there are other kinds of science, very different from physics, chemistry, geology and astronomy.

Science need not be reductionist

It is often assumed that a scientific explanation or description of any phenomenon must describe it in very low level terms, e.g. the language of physics referring to atoms, molecules, sub-atomic forces, etc.

But there are many examples of *organisation* of relatively simple things forming larger, more complex or more abstract entities for which new terms are needed: a “higher level” ontology.

- For example a chemist may refer to complex reactions involving large molecules or chains of reactions without having to reduce that to talk of sub-atomic particles and mechanisms.
- The changing patterns of motion of a large collection of fish may confuse predators. The high level changes can have real causal powers.
- Reproductive cycles in biology have important properties for explaining evolution or helping farmers breed new strains. We can talk about niches, species, genes, traits, capabilities, instead of talking only about physical particles and forces that make up the organism.
- Minds are not necessarily reducible to brains: different concepts are used for thinking about them: the ontologies are different, even if minds are implemented in brains.

Reality has many levels. Deep science will refer to many levels of reality.

Levels of reality

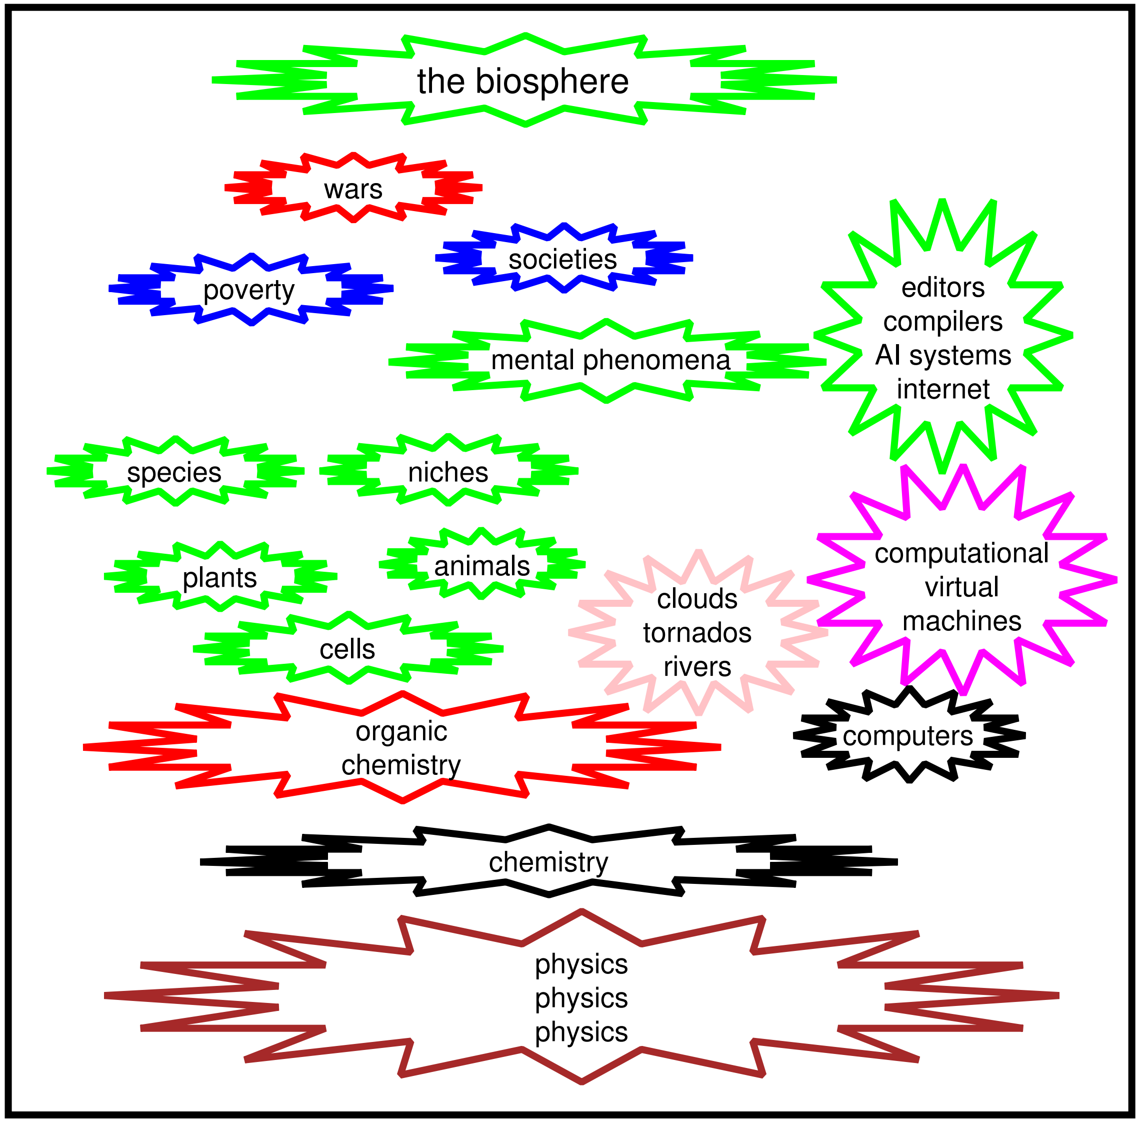

Physics has different levels: it seems to grow downwards over time.

Another kind of attempted reduction

Where do the concepts used to form a scientific ontology come from?

How can we decide which ones to use?

How are we able to understand them?

An old theory, which is often reinvented, is “Concept empiricism”, which claims that every concept an individual understands must either have been derived from experience of instances of the concept or defined by combining concepts that are so defined.

A recent version of this is known as “symbol-grounding” theory, presented by Stevan Harnad.
`http://cogsci.soton.ac.uk/~harnad/Papers/Harnad/harnad90.sgproblem.html`

Concept empiricism and symbol-grounding theory are both wrong

- In Critique of Pure Reason (1781) Kant argued that not all concepts could be based on individual experiences because experiences require concepts. So some concepts must have another origin.
- 20th Century philosophers of science proposed variants of concept empiricism in relation to scientific concepts e.g. by requiring every concept used to be “operationally definable”.

  This project failed, and eventually it was realised that some concepts, instead of being defined prior to being used in a theory, get their meanings primarily from their roles in theories.

  Why symbol-grounding and concept empiricism need to be rejected, and how the structure of a theory can provide partial definitions of concepts used in that theory, and how such a theory needs to be augmented with “bridging rules” that make the theory applicable, is explained in

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

  Why symbol-grounding is both impossible and unnecessary, and why theory-tethering is more powerful anyway.

Is there one truth to be sought by science?

A question often raised both by scientists and non-scientists is whether there are different kinds of theories favoured by different cultures, without any “objective” basis for choosing between them.

E.g. it may be suggested that there is some such difference between eastern and western science.

A partial answer:

- Any information-processing system (flea, frog, bird, human, robot) will use an ontology suited to its needs and capabilities.
- There is no sense in which the ontology used by a flea is wrong, if it suffices for flea-like purposes, and if the flea could not cope with a richer ontology.
- However, some organisms (including humans) include ontologies that can be (partly) mistaken or incomplete! (The flea’s ontology is incomplete, but it can probably never know that.)
- Some organisms (e.g. children, scientists) can investigate whether another ontology is better than the best they have found so far.
- This is how Newton’s ontology came to be replaced by Einstein’s in modern physics even though for many practical purposes Newton’s gives good approximations and is easier to use.
- However discovering that your ontology is incomplete or inferior to another can take a long time – and checking out which is better can be a very lengthy, expensive, and sophisticated process. There are no *simple* rules for choosing.
- Finding and removing gaps and errors is an important feature of human thinking.

Must science have practical goals?

A topic on which both philosophers and scientists are divided is whether science must have practical aims.

For example, some people think all scientific research must be justified by its potential economic value, or by the prospect of using the results for the benefit of humans.

However, alternative views are possible, e.g.

- Instead of aiming only to benefit human beings, science should aim to reduce the harm done by humans to other species, including those close to extinction – even if that harms or inconveniences humans.
- Science should aim to produce the best state of the universe, or at least the best possible state of our planet, irrespective of whether that involves preserving human beings or constraining their behaviours, or replacing them with something better – e.g. intelligent robots lacking the kinds of human nastiness involved in war, torture, murder, rape, religious fanaticism, racialism, nationalism, etc.
- The only core aim of science is to collect information about what is in the world and how things in the world behave in different conditions (using various forms of observation and experiment), and then to try to construct the best possible explanatory theory (or collection of theories) accounting for all the recorded observations.

  Proponents of this view will claim that if research done only for the sake of finding out how the world works has the side effect of helping humans, or helping other animals, that deserves commendation, but having that side effect is not a requirement for doing good science, and it need not be what motivates scientists.

Similar disagreements are possible about the aims of mathematics.

How to choose between the above positions is a topic for another time.

Science does not assume final answers are possible

Good scientists accept

- that no answer is ever final,
- that it is always possible that contrary evidence can turn up,
- that it is always possible that better theories will be suggested,

In this, science differs from many other types of activity, including most religious thinking, e.g. those which involve a commitment to *faith*.

This does not mean that science is a free-for-all, that “anything goes”.

If most people accept theory A as the best available in some branch of science, that does not mean that no scientist can propose theory B which is inconsistent with A.

It does mean that the reasons why B is better have to be articulated, and those reasons can then be investigated.

- It may turn out that the claim is spurious.
- It may turn out that B is far superior but only after new engineering technology has developed in order to reveal new evidence.
- It may turn out that deep analysis and testing does not determine which is better.
- In that case we have to go on developing the ideas until we see that they are actually equivalent, or we do find a reason for preferring one, perhaps in a hundred years’ time.

Three kinds of subject matter for science

Some scientists do not realise that the systems they study (e.g. humans and other animals, economic systems, social systems) are information-processing systems.

Science and engineering over the centuries have been concerned with many kinds of machines:

- Machines that manipulate matter

  E.g. levers, pumps, mechanical diggers, cranes, cookers.

  Of course they use energy.

- Machines that manipulate energy

  E.g. steam engines, dynamos, car engines, hydroelectric power stations, bicycles, lamps, batteries.

  This usually involves manipulating matter also.

- Machines that manipulate information i.e. meanings.

  Evolution got there long before human scientists did: all biological organisms use information in order to select between options for action, usually using stored chemical energy to produce the selected behaviour.

  The science of information is still in its infancy, and very few people understand how to think about it because it was not part of their education.

  (Using word-processors or internet browsers is not the same thing as learning the science of information, as some educationalists seem to think!)

What information-processing involves

So information processing involves: acquiring, storing, deriving, manipulating, inferring, analysing, interpreting, and above all using information in that familiar sense.

- This has no commitment as to
  - whether the processing is in machines or in organisms,
  - whether it is digital(discrete) or analog(continuous),
  - whether it is encoded explicitly and locally or implicitly in distributed form,
  - whether it is encoded physically (writing, pictures, ...) or in virtual machines
    e.g. in abstract data-structures, rules, axioms, networks, graphs,
  - Whether it is encoded within the information user or external to it, in the environment.
- All these are notions we need to analyse but don't have time for today.
  (Though every software engineer understands and uses them every day.)

See also:

Brian C Smith, 'The foundations of computing' in
M.Scheutz, ed., *Computationalism: New Directions*, MIT press, 2002.

The papers on varieties of representation here
`http://www.cs.bham.ac.uk/research/cogaff/`

What is information? 1

- We use “information” to mean:
  - Something like the ordinary notions of “content” and “meaning”
  - Not the Shannon-Weaver notion of information since:
    - information can be false.
    - items of information can stand in relations like consequence, contradiction and relevance
    - items of information can understood or misunderstood.
    - it has nothing to do with unexpectedness: information content is sometimes completely predictable

- Information is non-physical (albeit physically realised)
  - But this does not make it unsuitable for use in biology: compare “niche”, “gene”, etc.
  - It does, however, mean that specialised methodologies are required for identifying, and explaining, information processing.

    These differ from the methods of the physical sciences.
  - Compare the differences between:
    - identifying and fixing a faulty circuit in an electrical device, and
    - identifying and fixing a software bug – which may manifest itself in many physical implementations.

What is information? 2

The full answer is quite complex.

Partial answers can be found in talks 4 and talk 6 here:

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

And more detail here:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/whats-information.html

It is a mistake to seek an explicit definition of information

- Rather, “information”, like “energy”, will be implicitly defined by the role it plays in theories.
- Roughly, when you know:
  - the forms that information can take,
  - the variety of contents it can have,
  - the various ways it can be
    - acquired,
    - manipulated, analysed, interpreted,
    - stored, transmitted, tested,
    - and, above all, used,
- Then you know (to a first approximation) what information is.
- That knowledge evolves over time as we learn more, like our knowledge of what energy is.

What is information processing?

- In a sense we all know - so no answer is needed: it's what we refer to when we talk about
  - having information
  - acquiring information
  - needing information
  - using information,
  - truth, falsity, contradicting, being consistent, deriving, guessing ... etc.

- But we have only an intuitive understanding and there is a need to make this explicit, eventually: there's no time for a full analysis today.

  For more on this see

  e.g.
  `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/talk#4`
  also
  `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/whats-information.html`

- Clarifying "information processing" depends on clarifying the notion of "information", and that has several interpretations.

- For now we use only the ordinary notion of "information", linked to "meaning", "content", "reference", "inference", etc.

  Not the technical mathematical notion of information (Shannon's), which has little to do with meaning, but a lot to do with the science and technology of signal transmission.

NOTE: Some people assume information must be **true**. This is too restrictive – people can acquire, use, and communicate false information. There is also "control information", which is neither true nor false.

How to study information processing systems.

Regarding something as an information processor has important consequences.

This stance leads to very different research questions from a stance that views organisms as physical systems in a physical environment, to be studied using the methods of the physical sciences.

E.g. it provokes questions:

- What kinds of information does the system use?
- Where does the information come from?
- How is it acquired?
- In what forms is it expressed or encoded? (e.g. what syntax is used?)
- What is it used for?

  Motives, attitudes, preferences, desires can be regarded as “control information”, and emotions as “control states” - see other talks here.
- How is it stored, transformed, interpreted, used?

  E.g. in perceiving, learning, inferring, remembering, deciding, acting...
- What sort of architecture, using what sorts of mechanisms enable this?

These questions don’t make much sense if asked about a rock or a thundercloud – unless stretched a lot.

Is a science of mind possible?

Puzzling questions:

- Do minds exist, or are they some sort of illusion?
- How are they related to bodies?
- If brains do everything do we need minds?
- Is a science of mind possible? Useful? Urgently needed?

Sometimes you can get into a confused state because you are making assumptions you don't recognize

Then science needs to be combined with philosophy, as when Einstein asked

"what do we mean by two things happening at the same time in different locations?"

Compare some old 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)

Compared with the new information-processing ontology these ideas are mostly either circular or lacking in precise explanatory power.

Humans as information-processing systems

Trying to understand humans as information-processing systems, requires a rich and deep ontology for types of information, types of information storage, types of information mechanism, types of use of information (e.g. in motivating, deciding, acting, communicating, reasoning, ...).

We are not there yet, though we are making progress.

However:

- Some people have not even started: they still treat humans and other animals as if they were essentially physical/chemical systems to be observed, measured, and described using the language of the physical sciences, e.g. studying only physical and chemical properties of brains.
- Others abandon the language of physics but emulate a mistaken view of the methods of physics: they collect data and look for correlations.
- Others, who see that both of those approaches work badly, conclude that a scientific study of humans is impossible.

All three groups often reject the information-processing ontology because they don't understand what it is.

They wrongly believe that it is based on low-level comparisons between brains and computers.

What do information scientists study?

Many of them (perhaps most) study virtual machines, although the virtual machines need physical machines on which to run, and they are often connected to the rest of the world through physical devices, such as sensors and motors.

Virtual machines are abstract entities like:

- The processes inside a chess computer, analysing a board state, considering possible moves, using rules of chess to decide which options are legal, selecting moves on the basis of their likely consequences, etc.
- The processes inside a word-processor, spelling checker, email system, operating on text, fonts, formatting rules, words, sentences, paragraphs, etc.
- The processes inside an operating system like Windows or Linux, allocating other processes to memory, deciding which processes to run, keeping process statistics, checking file access privileges, connecting to processes on other computers, etc.

Note: “virtual” does not imply “unreal” in this context: virtual machines and the events that occur in them are real and can have real effects, including altering what’s on a computer display, or landing an airliner in a fog.

By creating many sorts of virtual machines over the last half century, we have begun to explore the variety of possible types. But we still have a very long way to go, especially if we want to understand the virtual machines in human minds and how they are related to the physical chemical machines we call brains.

Examples of research problems

Vision – perhaps the hardest problem in AI and psychology

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

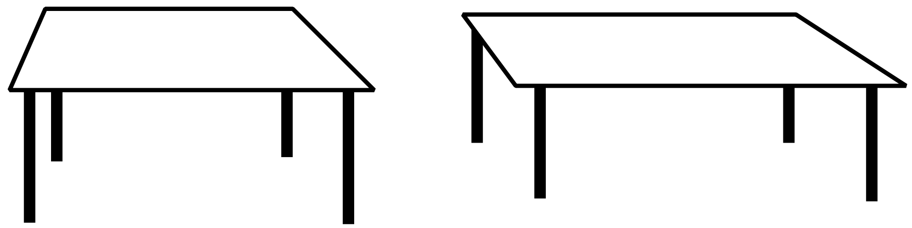

What if the objects are flexible, irregularly shaped, and moving?

How do we see expressions of emotion in faces?

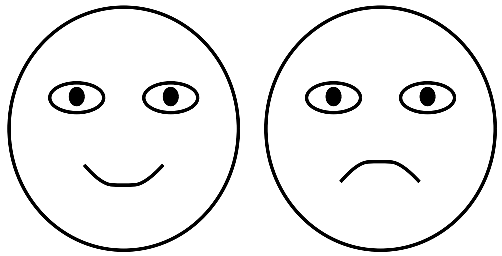

How are perceived emotions represented in a perceiver?

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

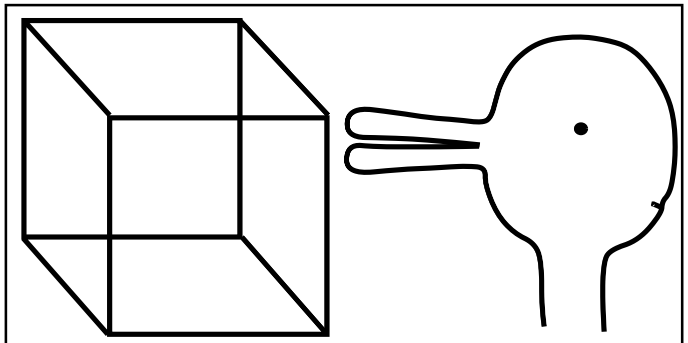

Necker cube Duck-rabbit

How are the differences represented?

And many more, including perception of motion, visual pleasure, study of motivation and emotions, learning to do mathematics, etc.

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.

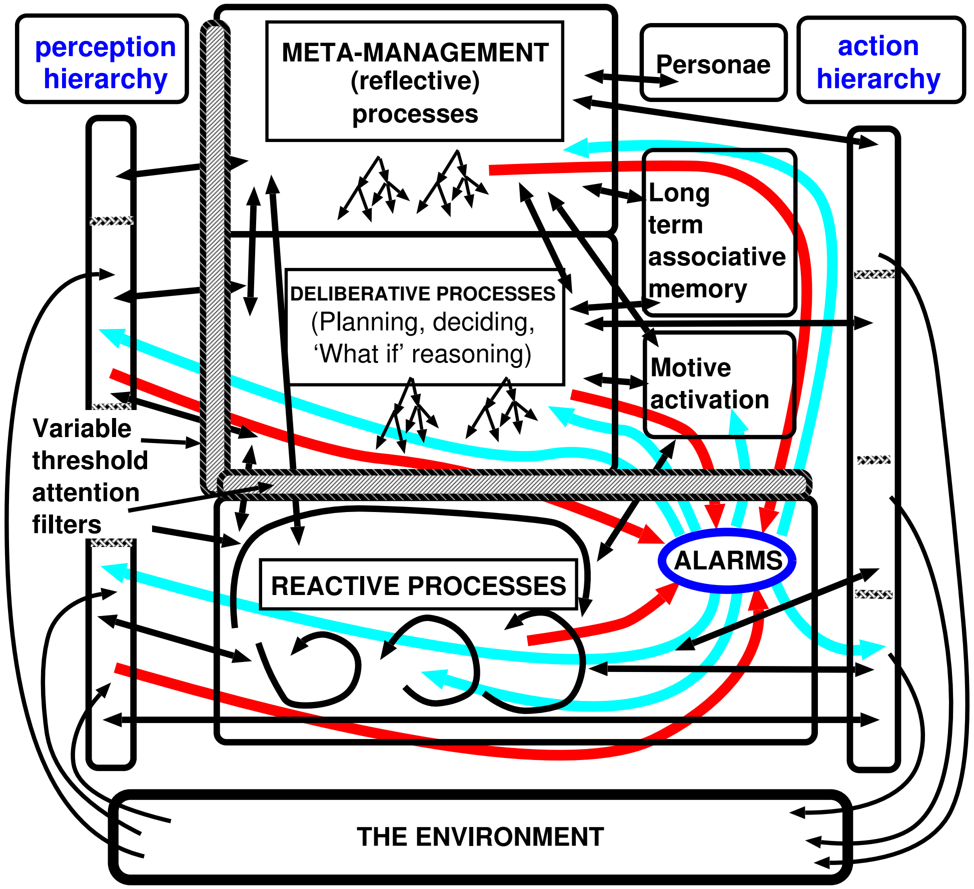

E.g. different sorts of emotions are generated in different levels.

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

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

This is too simple: A science of mind is possible, but very hard!

So, What then is science?

It is an example of a process in which information is acquired, manipulated and used. More precisely, if information users:

- constantly seek to find gaps and errors in their current best information about some part of the world
- constantly seek to find new ways to probe reality, e.g. using new instruments for making things happen and for collecting information
- constantly seek to modify and extend their concepts so as to allow more phenomena to be described and explained with greater accuracy and precision
- constantly seek new types of application of their theories
- constantly try to find widely applicable theories that are more precise, use fewer unanalysed concepts, have simpler but more powerful “axioms”, and explain and predict more distinct types of phenomena more precisely than in the past
- that can explain things the theory was not originally designed to explain
- preferably making some testable predictions that contradict widely held beliefs

Then they are doing science (and often engineering also).

Science is a collaborative and competitive questioning, self-criticising, self correcting, self-extending process that seeks out, acquires, transforms, integrates, and uses ever more detailed information, about ever more aspects of the world.

Both explicit testing and also using theories can show up new flaws to be addressed.

Overlaps and demarcation

Much of what I have said about science overlaps with and extends the common sense practices by which we attempt to understand the world, and constantly extend and improve our understanding – even when we are toddlers playing with toys – but there are differences.

Many attempts have been made to describe differences between science and other ways of trying to grow knowledge and understanding, e.g. in terms of

- the collection of evidence to support or test theories;
- the use of increasingly precise and sophisticated measuring devices;
- the use of mathematics in formulating theories and deriving consequences;
- assessing how probable it is that our theories are true, given the available evidence;
- requiring scientific theories to be falsifiable.

Falsifiability was elevated to a crucial test for science in Karl Popper's very inluential book The logic of scientific discovery, 1934.

Later, his pupil Imre Lakatos showed that a more subtle demarcation is needed, not between scientific theories and metaphysics or pseudo-science, but between progressive and degenerating research programmes.
`http://www.lse.ac.uk/collections/lakatos/scienceAndPseudoscienceTranscript.htm`

I tried to extend the ideas of Popper and Lakatos in Chapter 2 of The Computer Revolution in Philosophy (1978):`http://www.cs.bham.ac.uk/research/projects/cogaff/crp/chap2.html` showing how science is the study of what is and is not possible and why.

There's a lot more to be said about all this

Several slide presentations, exploring some of the issues in a science of information in more detail can be found here:

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

A web site with much useful information about Artificial Intelligence

`http://www.aaai.org/aitopics`

Some of the views on the nature of science expounded here are elaborated in this online book The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind (1978)

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

There is much more to be found in the writings of philosophers of science, including Karl Popper, Imre Lakatos, and others.

Try giving the string "what is science" to google and other search engines.

A critique of "Symbol Grounding Theory", the theory that all concepts have to be derived from experience (also known as "concept empiricism") can be found here

> `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models`
> Why symbol-grounding is both impossible and unnecessary, and why theory-tethering is more powerful anyway.
> (Introduction to key ideas of semantic models, implicit definitions and symbol tethering through theory tethering.)

Added January 2009:

> A web site reporting on a research project in Cognitive Robotics, funded by the European Union, is here
> `http://www.cognitivesystems.org`