How to make machines with souls

Souls as information-processing virtual machines

Aaron Sloman
http://www.cs.bham.ac.uk/˜axs/
The University of Birmingham

For related talks see

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

Hmmmm.......

How many people here have souls?
Does this computer have one?

A useful one, but not a human one.

THANKS

THIS MACHINE HAS A MICROSOFT-FREE SOUL

I am very grateful to
the developers of Linux
and other free, open-source,
platform-independent, software systems.

LaTex was used to produce these slides.

Diagrams are created using tgif, freely available from
http://bourbon.cs.umd.edu:8001/tgif/

Demos are built on Poplog

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

Is a soul a ghost in the machine?

What philosophers tend to forget - but the ghost of Gilbert Ryle knows well....

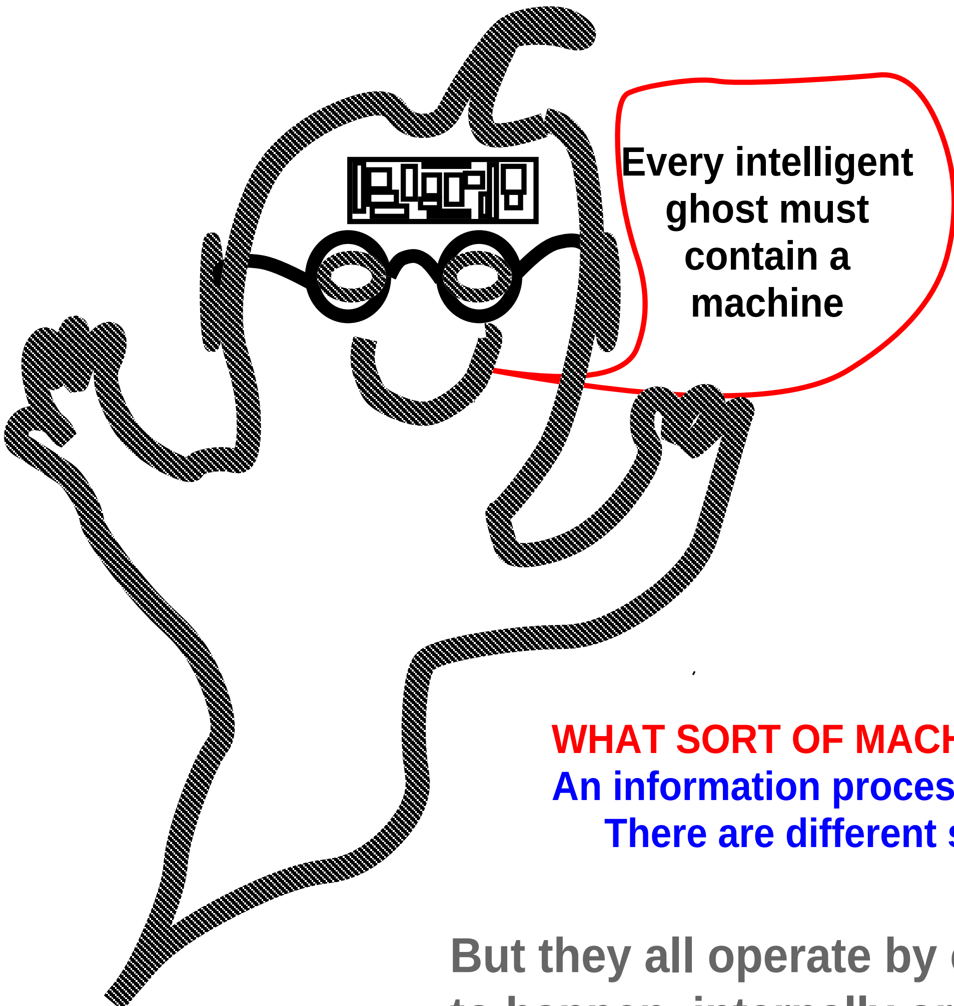

WHAT SORT OF MACHINE?

An information processing machine.
There are different sorts.

But they all operate by causing things to happen, internally or externally.

What sort of information processing machine?

SOME EXAMPLES
(time permitting)

Some simple examples of running virtual machines with different sorts of capabilities (different sorts of causal powers).

- **Talking machine (pretty dumb - does nothing unless sentence typed in)**
- **Moving-flocking**
- **Emoting**
- **Herding**
- **Interleaving reacting and planning**
- **Infant and carer**

Human-like souls will need complex architectures

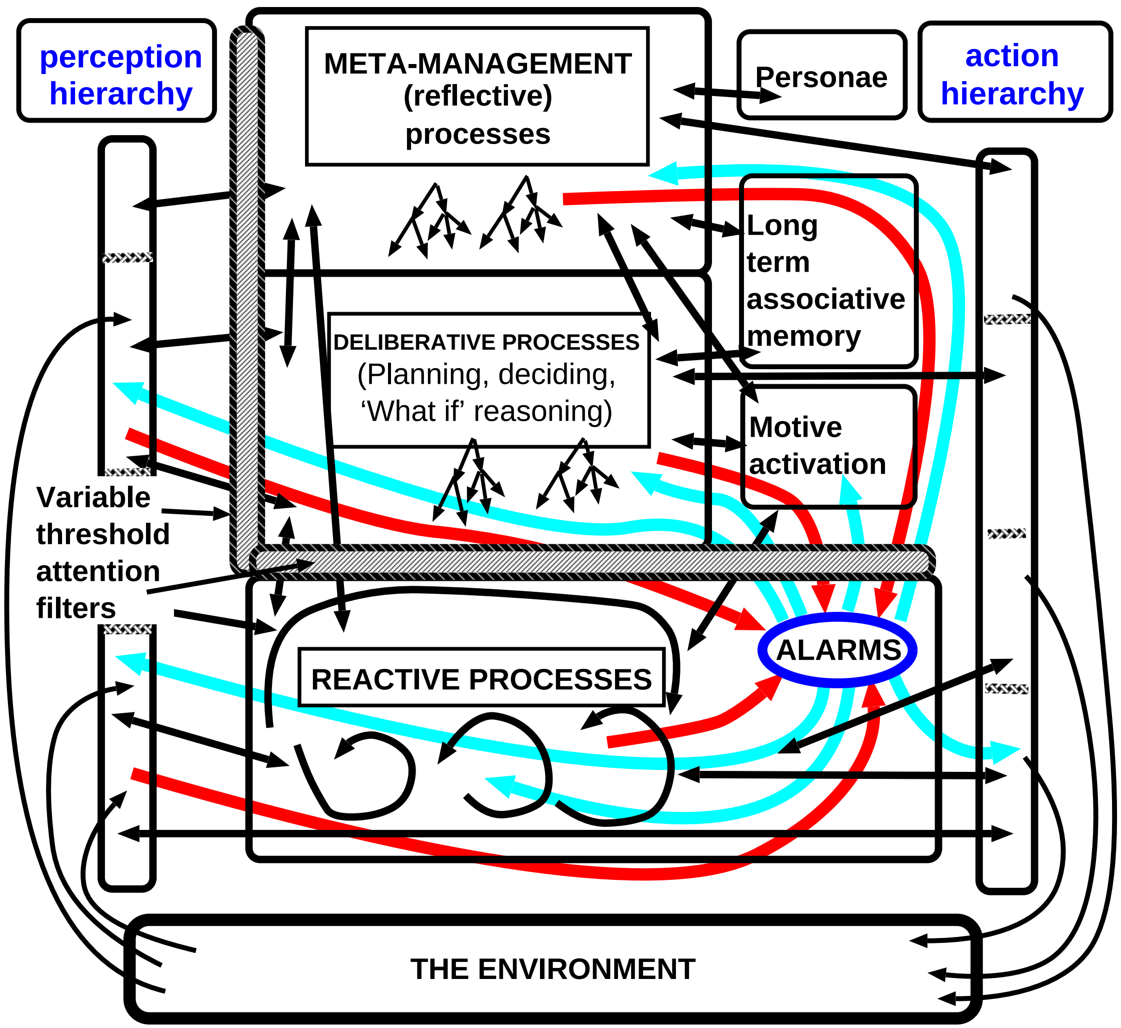

For more on this see the Birmingham Cognition and Affect project web site: http://www.cs.bham.ac.uk/˜axs/

Virtual machines are everywhere

At all levels there are objects, properties, relations, structures, mechanisms, states, events, processes and also many CAUSAL INTERACTIONS.

E.g. poverty can cause crime.

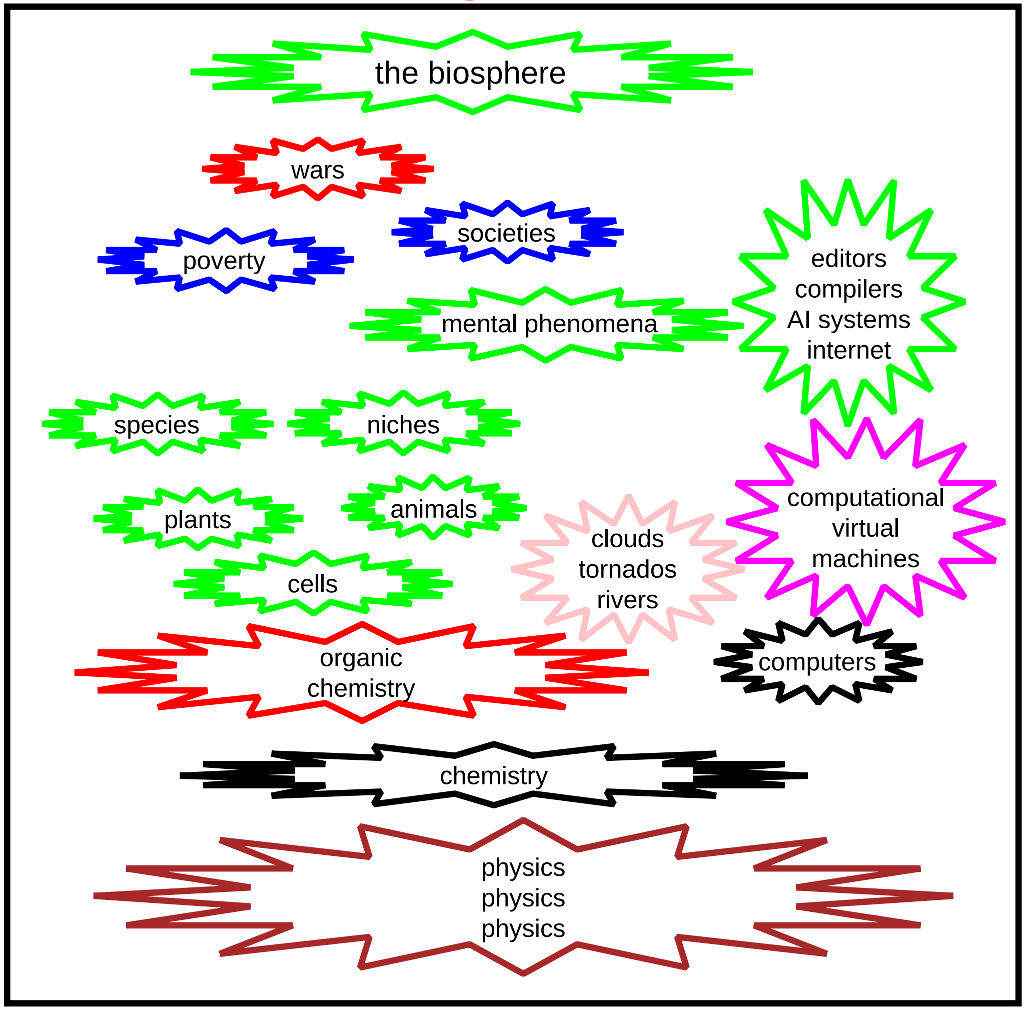

- **All levels ultimately realised (implemented) in physical systems.**
- **Different disciplines use different approaches (not always good ones).**
- **Nobody knows how many levels of virtual machines physicists will eventually discover. (Uncover?)**
- **Our emphasis on virtual machines is just a special case of the general need to describe and explain virtual machines in our world.**

See the IJCAI'01 Philosophy of AI tutorial (written with Matthias Scheutz) for more on levels and causation: http://www.cs.bham.ac.uk/˜axs/ijcai01/

Two notions of virtual machine

Some people object to claims

- **that causal interactions can occur within a virtual machine,**

and

- **that events in a virtual machine can be caused by or can cause physical events,**

because they ignore the difference between:

- **a VM which is an abstract mathematical object**
  (e.g. the Prolog VM, the Java VM, the Unix VM)

- **a VM that is a running instance of such a mathematical object, controlling events in a physical machine.**
  (E.g. the instance of linux running my machine now.)

The difference between these two is very important.

The mathematical object does not do anything (as numbers don't).

Running instances of virtual machines can do many things e.g.

- **landing a plane**
- **controlling a chemical plant**
- **monitoring patients in intensive care**

Anyone who claims that a virtual machine is just a formal entity has not understood these points.

Two notions of virtual machine

Contrast the notion of a PHYSICAL machine with:

- a VM which is an abstract mathematical object (e.g. the Prolog VM, the Java VM)
- a VM that is a running instance of such a mathematical object, controlling events in a physical machine, e.g. a running Prolog or Java VM.

Physical processes:	Running virtual machines:	Mathematical models:
currents	calculations	numbers
voltages	games	sets
state-changes	formatting	grammars
transducer events	proving	proofs
cpu events	parsing	Turing machines
memory events	planning	TM executions

VMs as mathematical objects are much studied in meta-mathematics and theoretical computer science. They are no more causally efficacious than numbers.

The main theorems, e.g. about computability, complexity, etc. are primarily about mathematical entities (and non-mathematical entities with the same structure – but no non-mathematical entity can be proved to have any mathematical properties).

There's more on varieties of virtual machines in later slides.

We need to extend our thinking capabilities and our ontologies

Many people are taught to think about

- Matter-manipulating machines
- Energy-manipulating machines

But they do not learn to think about

- Information-manipulating machines.

So they often fail to notice important questions and fail to consider important classes of possible answers: like neuroscientists who study neurons, and psychologists who study behaviour.

We need to extend our thinking capabilities and our ontologies

Many people are taught to think about

- **Matter-manipulating machines**
- **Energy-manipulating machines**

But they do not learn to think about

- **Information-manipulating machines.**

So they often fail to notice important questions and fail to consider important classes of possible answers: like neuroscientists who study neurons, and psychologists who study behaviour.

We are in the very early stages of learning to think about important age-old products of evolution:

- **Virtual machines:**
  - **with real causal powers**
  - **e.g. decisions change what happens.**
- **Much concurrency:**
  so that it can be misleading to ask what IT (or she or he) is doing, or can do, or notices, perceives, feels, etc.
  - **The answers may be different for different parts of the same system.**

Functionalism ?

Functionalism is one kind of attempt to understand the notion of virtual machine, in terms of states defined by a state-transition table.

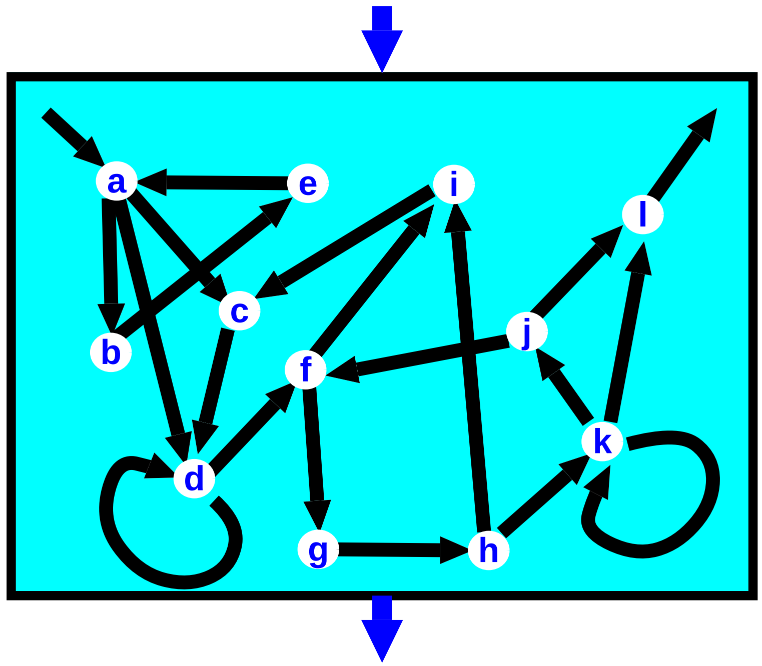

This is how many people think of functionalism: there's a total state which affects input/output contingencies, and each possible state can be defined by how inputs determine next state and outputs.

(E.g. see Ned Block's accounts of functionalism.)

HOWEVER THERE'S A RICHER, DEEPER NOTION OF FUNCTIONALISM

Another kind of Functionalism ?

Instead of a single (atomic) state which switches when some input is received, a virtual machine can include many sub-systems with their own states and state transitions going on concurrently, some of them providing inputs to others.

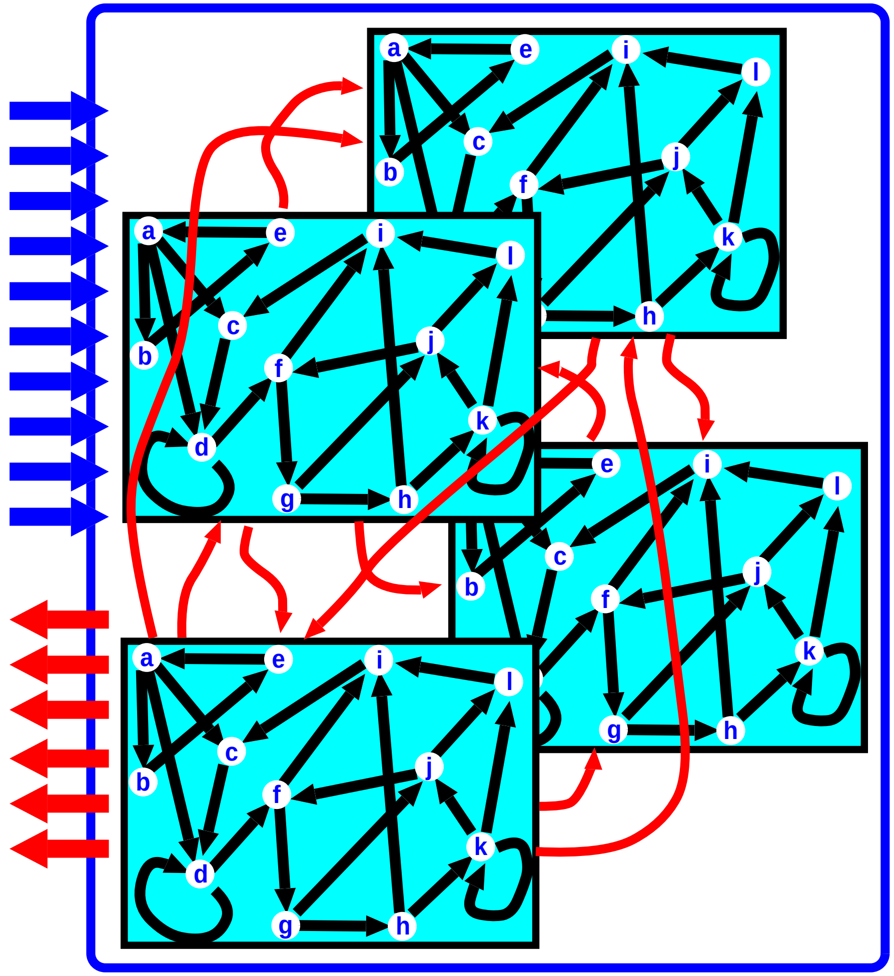

- The different states may change on different time scales: some change very rapidly others very slowly, if at all.
- They can vary in their granularity: some sub-systems may be able to be only in one of a few states, whereas others can switch between vast numbers of possible states (like a computer's virtual memory).
- Some may change continuously, others only in discrete steps.

Some sub-processes may be directly connected to sensors and effectors, whereas others have no direct connections to inputs and outputs and may only be affected very indirectly by sensors or affect motors only very indirectly (if at all!).

The previous picture is misleading

Because it suggests that the total state is made up of a fixed number of discretely varying sub-states:

We also need to allow systems that can grow structures whose complexity varies over time, as crudely indicated on the right,
e.g. trees, networks, algorithms, plans, thoughts, etc.

And systems that can change continuously, such as many physicists and control engineers have studied for many years, as crudely indicated on the bottom right.

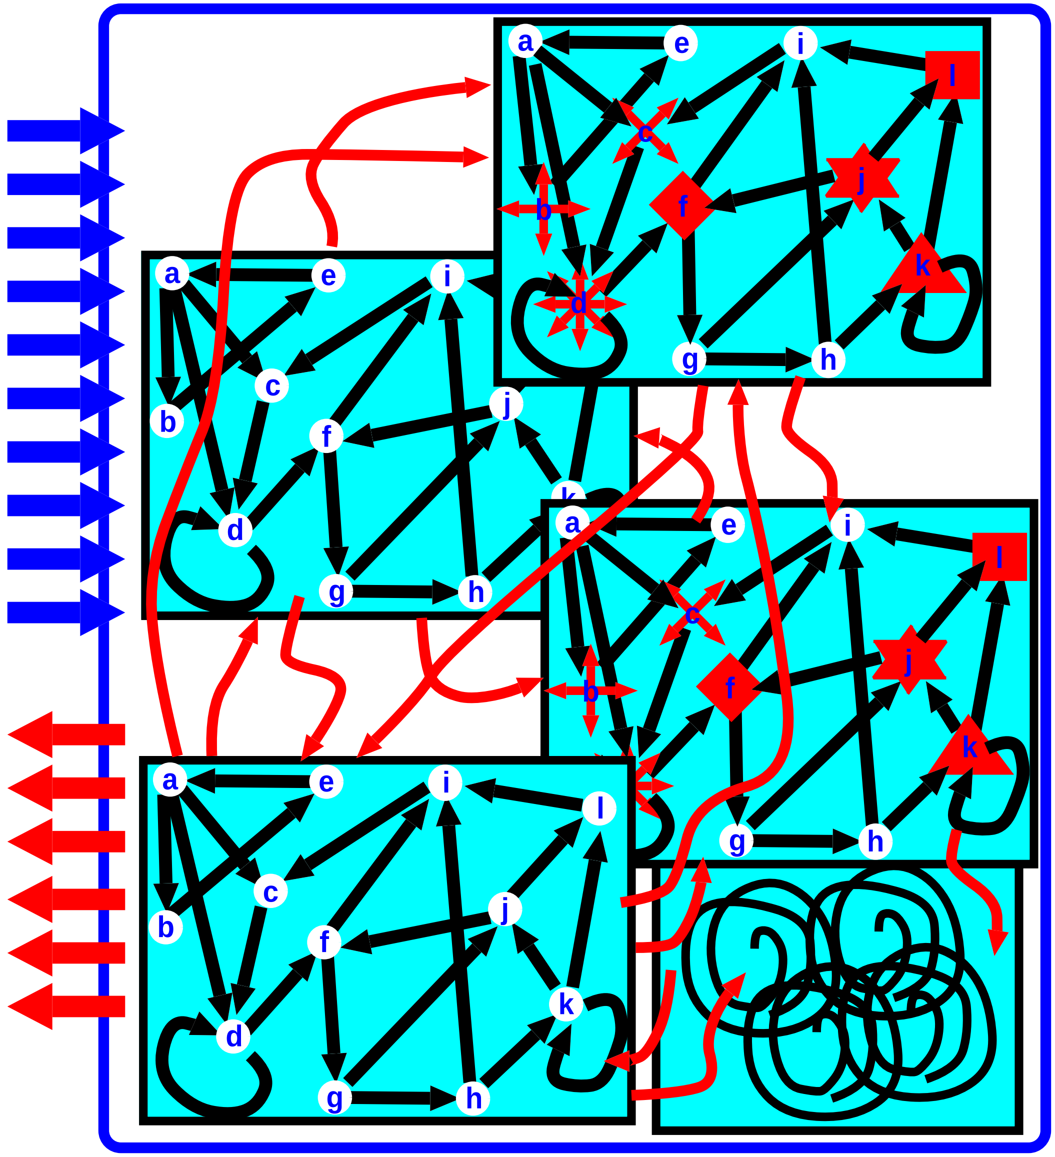

The label 'dynamical system' should be applicable to all of these types of sub-system and to complex systems composed of them.

VMF: Virtual Machine Functionalism

We use "Virtual Machine Functionalism" (VMF) to refer to the more general notion of functionalism, in contrast with "Atomic State Functionalism" (ASF) which is generally concerned with finite state machines that have only one state at a time.

VMF allows multiple concurrently active, interactive, sub-states changing on different time scales (some continously) with varying complexity.

VMF also allows that the Input/Output bandwidth of the system with multiple interacting internal states may be too low to reveal everything going on internally.

There may still be real, causally efficacious, internal virtual machine events and processes that cannot be directly observed and whose effects may not even be indirectly manifested externally.

> **Even opening up the system may not make it easy to observe the VM events and processes (decompiling can be too hard).**

If some links between systems can be turned on and off by internal processes, then during some states:

> **some of the subsystems may not have any causal influence on outputs.**

Those running sub-systems still exist and can include internal causal interactions within and between themselves: scientific investigations will have to allow for this possibility.

Get rid of the idea that a Turing test can be useful

The notion of a "Turing test" as something that can determine what is going on inside a complex system, fails to take account of many of the possibilities for virtual machines described on previous slides.

VMs can have temporarily or partly ‘decoupled’ components

- “Decoupled” subsystems may exist and process information, even though they have no connection with sensors or motors.
- For instance, a machine playing games of chess with itself, or investigating mathematical theorems, e.g. in number theory.
- It is also possible for internal VM processes to have a richness that cannot be expressed using the available bandwidth for motors.
- Likewise sensor data may merely introduce minor perturbations in what is a rich and complex ongoing internal process.

This transforms the requirements for rational discussion of some old philosophical problems about the relationship between mind and body:

E.g. some mental processes need have no behavioural manifestations, though they might, in principle, be detected using ‘decompiling’ techniques with non-invasive internal physical monitoring.
(This may be impossible in practice.)

Could de-coupled VM sub-systems be produced by evolution?

It is sometimes argued that sub-systems that do not have externally observable effects on behaviour would never be produced by evolution, because they provide no biological advantage.

This assumes an over-simplified view of evolution:

e.g. ignoring the fact that many neutral or harmless mutations can survive because they don't make sufficient difference to the survival chances of individuals. This could be because the environment is not sufficiently harsh or because more able individuals help less able ones or for other reasons.

A consequence is that a succession of changes that do not directly produce any great benefits (or disadvantages) may eventually combine to produce something very beneficial.

In some cases the benefits are insignificant until there's a major change in the environment requiring some new capability.

E.g. a succession of changes producing a mechanism for "thinking ahead" may be of no real benefit to members of a species until the environment changes so that food is not plentiful and actions to find food have to begin before the food is needed.

Likewise in individual development: virtual machines may change in (partly genetically programmed) ways that have no immediate benefit and show no behavioural consequences, but later on link up with other sub-systems and give the individual considerable advantages, e.g. mathematical thinking capabilities, perhaps.

The altricial-precocial spectrum

A full analysis of the issues requires discussion of the
Altricial-precocial spectrum

Architectures for altricial species can do many things that are not useful — though they could be.

Draft paper online

http://www.cs.bham.ac.uk/research/cogaff/altricial-precocial.pdf

Important point: Architectures can change over time.

Altricial architectures grow themselves.

We know very little about varieties of development and learning in virtual machines

- **Different models of development and learning are related to different starting points: Altricial/Precocial species (and machines).**
- **Precocial species have individuals almost completely determined by genes, whereas in altricial species there is a far more abstract genetic specification: a boot-strapping machine.**
- **Boot-strapping may be concerned with construction of a virtual machine, or virtual machine architecture, not just with wiring, etc.**
- **The fashion for ‘symbol-grounding’ theories of meaning ignores the richness of meaning that can be provided by internal structures and processes, e.g. results of millions of years of evolution.**
- **Kant: NOT all concepts can be learnt from experience.**
- **Many of the costs and constraints of biological systems are non-obvious: e.g. evolutionary history may or may not include opportunities for something to have evolved.**

Implications for testable theories

Virtual Machine Functionalism (VMF) implies that theories about systems using virtual machines can be very hard to test directly.

Instead we have to learn to work like physicists investigating sub-atomic entities, events and processes, where only very indirect testing is possible, and the most one can ever say of any theory is:

"This theory at present is better than any of its rivals"

It is always possible that a new, better, deeper, explanatory theory will turn up than we have discovered at any time, as happened when relativity and quantum mechanics replaced older theories.

This does not make truth relative, only very hard to discover.

Mental states and processes on this view are not mere "attributions" – they are real aspects of virtual machines.

Finding the right ontology for describing what's going on can be very hard: we still have much to learn about this.

Putting it all together

In the hope of reducing the confusion I have assembled these slides by collecting many partial explanations from papers and discussions over the last decade or so and modifying them in the light of what I've heard in recent debates. However:

- The issues are complex because the concepts used are not simple ones that can easily be defined explicitly.
- Moreover there are several different kinds of concepts involved, some relatively non-technical and widely understood, at least intuitively, others relatively technical and not well understood by most people.
- Some of the disputes depend on a view of computers that ignores the history that led up to them. For instance most of the key ideas were understood by Babbage and Lovelace long before the notions of Turing machine and equivalent mathematical notions had been thought of: computers are a recent development in a very old process of producing more and more sophisticated machines for controlling machines.
- Nowadays many of the controllers are virtual machines.
- It is also forgotten that computers were so-named because they were originally intended to take over a task that was previously done by humans, namely computing! (Likewise calculators performed a task previously done by humans.)

More importantly, living organisms have been processing information for millions of years.

Towards an ontology of 'mental' (i.e. VM) states

Our vocabulary for talking about virtual machines has two extremes:

- **the (very rich and powerful but very hard to analyse) concepts of ordinary language used when we talk about ourselves and other people**
- **the much more impoverished but much more precise and well understood concepts of virtual machines used in software engineering and AI, which are not yet adequate for characterising biological systems.**

We need to move towards something in-between, which is both precise and relevant both to organisms and machines, e.g. states that classified in The Architectural Basis of Affective States and Processes as:

- **Belief-like**
- **Desire-like**
- **Supposition-like**
- **Plan-like**
- **Moods and other varieties of affect**
- **initiation, termination, modulation, arbitration, evaluation ...**
- **Emotions as perturbances of one part by another**

We can see the required variety of types of VM states by considering diverse biological organisms, from microbes to elephants.

A biological perspective

Once upon a time there were only inorganic things:
atoms, molecules, rocks, planets, stars, etc.
These merely reacted to *resultants* of all the physical forces acting on them.

Later, there were simple organisms. And then more and more complex organisms.

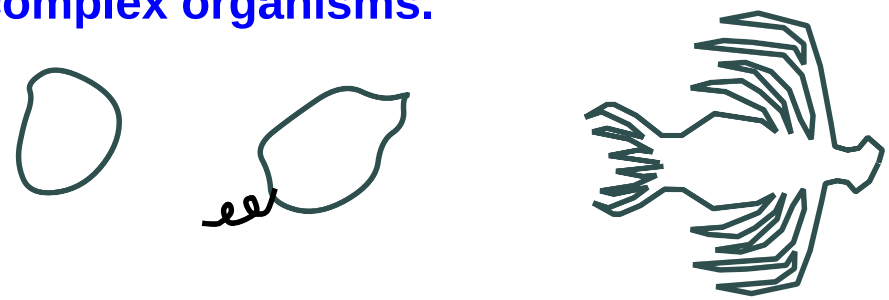

These organisms had the ability to reproduce. But more interesting was their ability to *initiate* action, and to *select* responses, instead of simply being pushed around by resultants.

That achievement required the ability to acquire, process, and use *information.*

The ability to act or to select requires information

E.g. organisms can use information about

- density gradients of nutrients in the primaeval soup
- the presence of noxious entities
- where the gap is in a barrier
- precise locations of branches in a tree as you fly through
- how much of your nest you have built so far
- which part should be extended next
- where the nest is, or where a potential mate is
- something that might eat you
- the grass on the other side of the hill
- what another animal is likely to do next
- how to achieve or avoid various states
- how you thought about two problems, one solved the other not
- whether your thinking is making progress ... and much, much more...

All this requires that organisms contain an energy store which can be deployed to meet their requirements, unlike most physical objects whose behaviour is determined only by external forces.

In a bouncing ball, elastic energy is temporarily stored, put there by physical forces, then released in a manner that has nothing to do with a need for survival of the ball. The ball uses no information: it has no needs or purposes — It takes no steps to survive or reproduce.

The notion of need

- **Making all that precise requires the notion of a need and a process or mechanism that serves the need.**
- **The existence of such things amounts to the truth of very complex sets of counterfactual conditional statements**
  - **About what would or would not happen in various circumstances if the need were not satisfied.**
  - **About what would or would not happein various circumstances if the need-serving process or mechanism did not exist or were modified in some way.**

The evolution of information-processing

Over time, as organisms became more complex, their use of information became more complex.

- Instead of reacting immediately to sensed states and events, some evolved the ability to take in information and use it later, e.g. going back to a location where food had been perceived.
- Some evolved the ability to make their reactions to particular sensed stimuli depend on internally sensed states of need.
- Some evolved the ability to allow more than one reaction to be triggered simultaneously and to use sensed or stored information influence the choice when the reactions are incompatible.
- Some evolved the ability to react to derived information, e.g. inferring the presence of a predator nearby and reacting to the derived information.
- Some developed the ability to acquire, store and use, possibly much later, generalisations about things in the world.
- Some developed the additional ability to derive and compare two or more predictions or plans, compare them and then select one. This required means of encoding hypotheticals.
- Some developed the ability to acquire and use information about their own information-processing, or information about the information-processing done by other individuals, e.g. predators, prey and neutral individuals

Some qualitative changes

Many assume biological evolution is a continuous process: but it cannot be (a) because DNA cannot change continuously - molecules are discrete structures, and (b) because there are only a finite number of generations between any two states.

- **One of the important qualitative changes involved being able to discretise or chunk information: this is necessary to explore branching sets of possibilities, whether for exploring alternative sequences of action in making a plan, or exploring alternative sequences of other kinds in making predictions, or exploring alternative explanations for observed facts.**
- **That change led to requirements for new processes of perception, new forms of information storage, new kinds of temporary work-spaces, new ways of managing decisions.**
- **Another kind of qualitative change was development of means of acquiring and using information about the activities of an information user, whether oneself or another individual. This required an extension of the ontology beyond what was adequate for expressing information about physical objects and their interactions in the environment.**
- **We still do not know enough the requirements for these changes, nor about the possible kinds of mechanisms that can support them, nor which kinds of architectures can combine these and other kinds of information-processing. (But we know much more than we knew a hundred years ago.)**

Varieties of biological information-processing

Different animals (microbes, insects, fishes, reptiles, birds, mammals, etc.) clearly differ in their requirements and their capabilities.

It would be helpful to attempt a survey of "dimensions" in which such capabilities can vary, and the kinds of designs that can support the different varieties.

This would be part of a general theory of information – what it is and how it works.

One of the kinds of dimensions would be concerned with the sort of *content* of the information.

- **Some information is very localised and simple (here's a dot, there's some motion to the left).**
- **Other information is far more holistic (e.g. recognising a scene as involving a forest glade).**
- **Some may be very abstract (the weather looks fine; it looks as if a fight is about to break out in that crowd).**
- **Some information items contain generally applicable knowledge, e.g. about the geometry and topology of static and moving shapes: e.g. regular hexagons can be packed to fill a convex space.**
- **Others involve specific facts relevant only in a particular part of the world, e.g. the Eiffel tower is in Paris.**
- **Some items of information are "categorical" others "hypothetical" or counterfactual, e.g. you would have been killed by that car had you not jumped out of its way.**

Other modes of variation are concerned with the medium used and the formal or syntactic properties of the medium.

Capabilities of different organisms and different machines

Some steps required for a more complete theory.

- If we develop a good ontology for types of information contents, we can start asking which organisms can handle which kinds.
- It is not clear which kinds of information contents different animals are capable of creating, understanding or using, or why: this presumably is related to their mechanisms, forms of representation and architectures.
- Likewise it is not clear which kinds children can cope with at different stages of development.
- A good theory would help us explain why certain types of robots are, and others are not, capable of acquiring, understanding, using certain sorts of information.

Can we do all this work without first defining “information” ?

What is information?

Resist the urge to ask for a DEFINITION of "information"

Compare "energy" – the concept has grown much since the time of Newton. Did he understand what energy is?

Instead of *defining* "information" we need to analyse kinds of processes in which it can be involved, the kinds of effects it can have, and the kinds of mechanisms required, i.e. such things as

– the variety of types of information there are,
– the kinds of forms they can take.
– the variety of means of acquiring information,
– the means of manipulating information,
– the means of storing or transmitting information,
– the means of communicating information,
– the purposes for which information can be used,
– the variety of ways of using information.

Examples of all of these will be given later

As we learn more about such things, our concept of "information" grows deeper and richer: Like many deep concepts in science (including "energy" and "matter"), the concept of "information" is mostly *implicitly* defined by its role in our theories and our designs for working systems.

Compare “information” and “energy”

It is also hard to define “energy” in a completely general way.

Did Newton understand the concept “energy”?

There are kinds of energy he did not know about:

- **chemical energy**
- **electromagnetic energy, ... etc.**

Why were these called “energy”? The theory that energy is *conserved* was crucial.

We can best think of energy in terms of:

- **the different forms it can take,**
- **the ways in which it can be**
  - **acquired**
  - **transformed,**
  - **stored,**
  - **transmitted,**
  - **used, etc.**
- **the kinds of causes and effects that energy transformations have,**
- **the many different kinds of machines that can manipulate energy**
- **....**

If we understand all that, then we don’t need to *define* “energy” – at least not by specifying its meaning in tems of ways of testing or measuring the presence of energy.

It is a primitive theoretical term – implicitly defined by the processes, relationships and mechanisms that involve it.

How not to define deep theoretical concepts

Newton knew about energy, but did not know anything about the energy in mass:

The possibility of $E = MC^2$ had not been thought of.

(This partially transformed both the concepts "energy" and "mass".)

We should not use currently known forms of energy or current ways of measuring energy to *define* it, since new forms of energy may turn up in future, along with new types of measurement.

(Partial changes to the theory partially change the concepts.)

This is typical of deep scientific concepts: they are to a large extent implicitly defined by the theories in which they are used, and cannot be explicitly defined in terms of pre-theoretical concepts or types of measurements or observations.

Any such definitions ("operational definitions") would omit central features of the concepts, namely their structural and causal connections within the theory.

Note: All this is familiar to philosophers of science, but not always understood by scientists, especially those who think physics and chemistry are merely about laws relating observables.

A related confusion is the wide-spread "concept empiricist" belief that all concepts must somehow be abstracted from experience, sometimes labelled the theory of "symbol grounding". Concept empiricism (and therefore symbol grounding theory) was demolished long ago by Immanuel Kant. See http://www.cs.bham.ac.uk/research/cogaff/talks/#talk14

"Getting meaning off the ground: symbol grounding vs symbol attachment"

Contrast Shannon's notion of "information"

We are not using Shannon's syntactic notion of "information" which refers to statistical properties of possible collections of symbols.

We are using something closer to the colloquial notion of "information" as

- **meaning**
- **reference**
- **semantic content**

which requires there to be

1. **a user or interpreter of the meaning (recipient, in the case of a message)**
2. **a bearer, or encoding, of the meaning (a picture, sentence, dance, wave pattern, electronic state of a memory chip, etc.)**
3. **sometimes, but not always, there is a source of the encoding (e.g. sender of a message)** **(Source, or creator, and recipient or user, are often one thing.)**
4. **something which is expressed or referred to (the content)** **(Mill, Frege and others distinguished two aspects: sense/connotation/intension and reference/denotation/extension)**

For more on this see: http://www.cs.bham.ac.uk/research/cogaff/talks/#talk14

Note:

Some "information-bearers" are physical (e.g. marks on paper), but often the bearer is a structure or process in a virtual machine. E.g. a network data-structure in a computational virtual machine could encode, for that machine, information about a network of roads, used by a route-finder.

Differences between energy and information

We are not using a *quantitative* notion of information

One big difference between energy and information (in the sense used here):

It is very useful to *measure* energy e.g. because it is conserved.

Expressing information as a numerical quantity is often of no use.

Numbers describing information (measurements) are *sometimes* useful

(e.g. if one message contains information about three people and another contains additional information about a fourth person).

But numbers do not capture what is most important about information, for behaving systems:

Numbers don't express where something is (e.g. in a drawer), what it is, how it is related to other things, where it comes from, what it can do, who made it, what the implications of something are, etc.

Further differences:

- **If I give you information I may still have it, unlike energy.**
- **You can derive new information from old, and still have both, unlike energy.**
- **Information varies primarily not in its *amount*, like energy, but in its structure and content: numeric equations do not represent most information manipulations adequately.**
  (Compare chemical equations, parse trees, maps, flow-charts.)
- **Energy in a physical object is there independently of whether any machine or organism takes account of it, whereas the information expressed or conveyed by something depends on the information-processing capabilities of the user or perceiver: information (in the sense we are using) is inherently relational.**

Being relational does not imply being subjective

- Whether a jacket J is a good fit depends on who the wearer is.
  So being a good fit is a relational property.
- But if X is a particular person, then whether J is a good fit for X is not a relational property.
- Neither is it merely something arbitrarily attributed to J by perceivers.

- Likewise what information a particular information-bearer expresses will depend on who is attending to the information.
- However potential information content for different sorts of perceivers is an objective property: so

  the statement that an object O can convey information I to agents with certain kinds of information processing capabilities C is not just an arbitrary or subjective attribution: it's a fact about the relationship between features of O and I and C

- Checking its truth may be very difficult however.

Things that can be done with information

Part of an analysis of the notion of “information” is provided by a taxonomy of types of things that can be done with information, by a user or perceiver X:

- **X can react immediately (the information can trigger immediate action, external or internal)**
- **X can do segmenting, clustering, labelling of components within a complex information structure (i.e. do parsing)**
- **X can try to derive new information from something (e.g. what caused this? what else is there? what might happen next? can I benefit from this?)**
- **X can store the information for future use (and possibly modify it later).**
- **X can use the information in considering alternative next events, in making predictions.**
- **X can use information in considering alternative next actions, in making plans**
- **If X interprets some information as containing instructions, X can obey them, e.g. carrying out a plan.**
- **The information can express one or more of X’s goals, preferences, ideals, attitudes, etc.**
- **X can observe itself doing some or all of the above and derive new information from that (self-monitoring, meta-management).**
- **X can communicate the information to others (or to itself later)**
- **X can check information for consistency, either internal or external**
- **X can check information for correctness (truth), precision, relevance, ...**

and lots more .....
using different forms of representation for different purposes.

Sentences, lists, arrays, metrical maps, topological maps, pictures, 3-D working models, weights in a neural net, structures of complex molecules, data structures in a computer, gestures, etc.

Diverse mechanisms of varying sophistication

Extracting information from basic sensory data may require very different perceptual mechanisms with varying sophistication.

- Some information can be extracted very simply (using spatial or temporal local change detectors, or mechanisms for constructing histograms of features, such as colour, texture, optic flow).
- Other information may need *relationships* to be discovered between features, e.g. collinearity, lying on a circular arc, parallelism, closure, lying on the intersection of the continuations of two linear segments or two curved segments (where the continuations are also curved).
- Sometimes this requires *searching* for coherent interpretations.
- Some relationships hold only between abstract entities not the image data: e.g. two people seen to be *looking in the same direction*.
- Extracting some of the information requires matching with known models ("That's a triangle, a face, a tree").
- Some learning tasks require noticing new repeated structures within the information structures (e.g. noticing repeated occurrence of polygons with circles at two adjacent corners).

For different kinds of sensory interpretation tasks, different forms of representation are often useful, and different types of processing.

There are different kinds of information

For instance:

- **about categories of things (big, small, red, blue, prey, predator)**
- **about generalisations (big things are harder to pick up)**
- **about particular things (that thing is heavy)**
- **about priorities (it is better to X than to Y)**
- **about what to do (run! fight! freeze! look! attend! decide now!)**
- **about how to do things (find a tree, jump onto it, climb...)**

This categorisation of types of information does not cover all the types found in machines and organisms.

Some of the differences are differences in "pragmatic function" rather than "semantic content".

We probably still know only about a small subset of types of information, types of encoding, and types of uses of information.

Don't expect all types to be expressible in languages we can understand – e.g. what a fly sees, or what a bee expresses in a dance!

Or even what a chimp, or a human child sees

We often tend to ask whether an animal can learn that so and so without considering the the implications of the possibility that nothing the animal is capable of learning is expressible in a human language or thinkable in a human mental architecture.

Further aspects of a theory of information.

We need to understand other ways in which information-processing events can vary.

E.g. besides

- **Different information contents, and**
- **the different forms in which they can be expressed,**

there are further functional and causal differences:

- **the different ways information can be acquired, transformed, stored, searched, transmitted, combined or used,**
- **the kinds of causes that produce events involving information,**
- **the kinds of effects information manipulation can have,**
- **the many different kinds of machines that can manipulate information,**

If we understand all that, then we don't need to *define* "information"!

CONCLUSION

Souls are of interest only insofar as they sense and act, on the environment, and on themselves.

We don’t yet know what souls are because we don’t yet know all the things they do.

As we learn more and more about what they do, so will we increase our chances of knowing how they can be designed and built.

It is very likely that some aspects can be designed and others will have to be built by the working souls themselves as they discover what sort of environment they are in and how they need to interact with it.

Please join the project:
only 300, or maybe 3000, more years to go.

THANK YOU