What Has Life Got To Do With Mind?

Or vice versa?

(Thoughts inspired by discussions with Margaret Boden.)

See:

M. A. Boden, Autopoesis and life, in *Cognitive Science Quarterly*, 2000, 1, 1, pp. 115–143,

M. A. Boden, 2006, *Mind As Machine: A history of Cognitive Science* (Vols 1–2), OUP 2006
Chapters 15.x.b and 16.x

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

These slides will go into my ‘talks’ directory:
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#maggie

Schedule for the Meeting

Friday 22nd May 2009, Arts Building B
University of Sussex, Falmer

14:00-14:45 Aaron Sloman

14:45-15:30 Ron Chrisley

15:30-16:00 discussion 1

16:00-16:30 break

16:30-17:15 Mike Wheeler

17:15-18:00 Steve Torrance & Blay Whitby

18:00-18:30 discussion 2

Chairperson: Paul Davies

Organiser: Tanja Staehler

Margaret Boden

Some of the information on her web site

Fellow (and former Vice-President) of the British Academy – and Chairman of their Philosophy Section until July 2002.

Member of the Academia Europaea.

Fellow of the American Association for Artificial Intelligence (AAAI).

Fellow of the European Coordinating Committee for Artificial Intelligence (ECCAI).

Life Fellow of the UK's Society for Artificial Intelligence and the Simulation of Behaviour (SSAISB).

Member of Council of the Royal Institute of Philosophy.

Former Vice-President (and Chairman of Council) of the Royal Institution of Great Britain.

OBE in 2001 ("for services to cognitive science")

Cambridge ScD

Three honorary doctorates (Bristol, Sussex, and Open University).

A few random reminiscences

- I met Maggie at one or more philosophy of science conferences somewhere early 1960s.
- I came to Sussex U. in 1964. Maggie came from Birmingham (as Phil&Psy lecturer) 1965: a great catch.
- I have one regret about that time .... (explained later)
- She was one of the group that started up a 'Communication Studies' programme in school of Arts and Social sciences, in 1973. (Why not 'Cognitive'? Stuart Sutherland had claimed that.)
  We also provided an AI programming prelim for students in the Arts and Soc Area.
  Alistair Chalmers, Max Clowes (transferred from EP), John Lyons, ...
  Supported by the Dean: Donald Winch
  Later we appointed Steve Hardy, Gerald Gazdar, Chris Mellish, Jonathan Cunningham, ...
  others were interested, including Barbara Lloyd for a while.
  Julie Rutkowska and later on Stephanie Thornton, in Developmental Psychology. and Mike Scaife in Social Psychology.
  We collaborated with people in Exp Psy, e.g. Christopher Longuet-Higgins, Steve Isard.
- Arts E, PDP11/40, terminals, and debugging room, 1975: Steve Hardy, Pop-11, Rhino, Eliza, ...
- Lots of team teaching. Tried programming tutorial for non AI staff.
  (Whoops of delight from Maggie when her programme worked – unlike some other people).
- We introduced an MSc conversion course around 1984-5 thanks to the UK Alvey Programme, with enormous help from Margaret McGowan (professor of French and ProVC Arts and Soc!)
- Peter Lloyd, then Dean suggested we have an AI undergraduate major (the first in the world?)
- Funded partly by selling Poplog (massive design/development contribution from John Gibson).
- Soon after that we joined with computer science (then in MAPS, led by Matthew Hennessy), formed a new School, COGS, and Maggie was unanimously approved as first Dean. Soon moved over to building between MAPS and Chemistry. Had to reinvent the "debugging room".
- I left in 1991 after 27 years, to her old university, but continued interacting with Maggie
  (Ved support, conferences, email).

An overview of her work could take up several seminars

Amazing achievements, starting with articles she was writing in the 1960s

(I didn't know till after 1969: to my regret: so it was Max Clowes who first taught me about AI.)

FIRST BOOK **Purposive Explanation in Psychology** 1972 (Harvard PhD bridging 3 disciplines)
(First ever non-experimental PhD in psychology at Harvard: contained seeds of much of her later work.)

Artificial Intelligence and Natural Man, 1978 ("footnotes to her first book".)
The first **general** intro to AI/Cognitive Science + philosophy of both. Very high impact. E.g. OU textbook.
(She also made me write a book, one of several published (1978) by her new company Harvester Press.)
She was awarded many honours, See earlier slide. OBE announcement:

Professor Margaret Boden, an internationally respected authority in the field of artificial intelligence and cognitive science, has been awarded an OBE for services to cognitive science in the New Years Honours List.

Many other papers, books, invited presentations, ...

MAGNUM OPUS: Mind as Machine: A History of Cognitive Science OUP June 2006
1,631 pages (small font!) in two volumes, one purple one pink,

An astounding achievement. Much, much more than a history. Will be influential for decades, or more...

I can't wait to get a digitised version, to keep on my computer
(to read, to browse, to search... with large font!)

The list of honours and invitations will grow – e.g.

2009 Turing Memorial Lecture

TURING & ARTIFICIAL LIFE - An Evening Lecture by Professor Margaret A Boden
Hosted by: Bletchley Park Trust 02 July from 18:00 to 22:00: lecture then dinner.

There will be many more to come....

What we have in common: Maggie started first

We are both engaged in work that combines philosophy, science and engineering and will, eventually, transform all three.

Describing the full details would take too long, so here's a sketch of some of them:

Key ideas:

- Doing philosophy in general, and conceptual analysis in particular, in ignorance of new discoveries and achievements in science and technology, is a recipe for a sickly inbred species of thinking.
- New advances outside philosophy provide a new platform on which to build conceptual investigations.
- The outcome of conceptual analysis, called logical geography by Gilbert Ryle, can be seen as a transient/culturally-based patchwork imposed on an enduring terrain.
- The underlying logical topography is discovered (gradually) by non-philosophical (scientific and technical) advances.

  Compare

  (a) the early theories of kinds of stuff (earth, air, water, wood, carbon, stone, etc. – the old logical geography) with

  (b) what came after discovery of the architecture of matter and the periodic table of the elements, plus chemistry (new logical topography, spawning new logical geographies).

Examples of risks to philosophy

- At one point a philosopher analysing common concepts would have included whales, dolphins, etc. as types of fish, e.g. because of their form, habitat and behaviours.

  As a result of evolutionary theory and discovery of more empirical facts they are now classified as mammals, requiring a revision of our concepts of types of animal.

- Philosophers of mind ignore biology, psychology, neuroscience and developments in computation at their peril.

- Philosophy of science, which used to pay attention mainly to physics, needs to look at new sciences, including computer science, AI, genetics, linguistics, ...

  (Of course, some philosophers are already doing this)

- Philosophical attempts to understand causation that ignore causation in virtual machines, are out of date!

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

See Also

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0703`
Two Notions Contrasted: ‘Logical Geography’ and ‘Logical Topography’
(Variations on a theme by Gilbert Ryle: The logical topography of ‘Logical Geography’.

Two Types of Logical Ography

Logical geography – description of a network of currently used concepts :

Often assumed by linguistic/analytical philosophers to be a kind of non-empirical truth – not just empirical facts about usage in some community.

Example:

A conceptual analysis claiming that All bachelors are unmarried is a conceptual truth, is close to this empirical claim:

The word "bachelor" (in one of its sense, among English speakers) is used to mean "unmarried adult male..."

Compare: The paradox of analysis

(My old tutor, R.M.Hare and others struggled with this.)

Logical topography – description of the structure of some aspect of reality, supporting a space of possible geographies,

as revealed by science and technology as well as common observation.

A new role for philosophy: helping the sciences clarify the logical topographies they uncover, and investigating pros and cons of alternative logical geographies based on those logical topographies.

(Not necessarily normative: the aim is explanation not instruction.)

Logical topographies underpin new logical geographies

A logical geography carves up some region of reality in a particular way, for particular purposes, partly under the influence of a particular cultural history.

E.g. an ontology for kinds of matter and kinds of transformations of matter

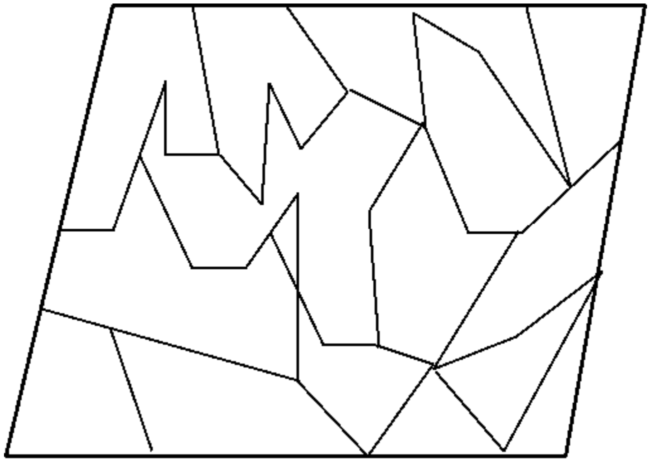

The advance of science and technology can gradually reveal a richer underlying reality – the logical topography – that can be divided up in various different ways.

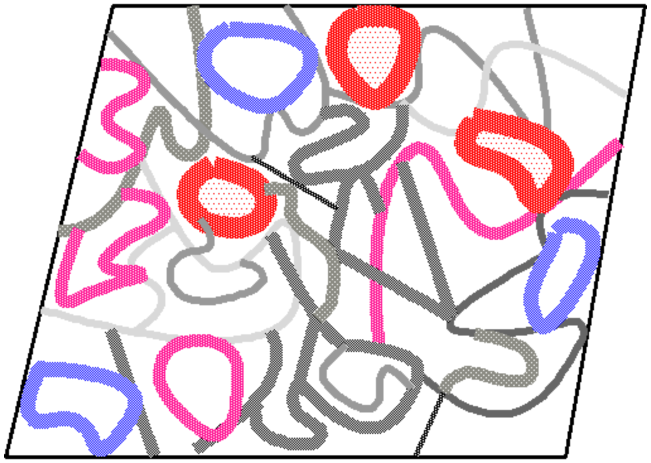

The process of "uncovering" the underlying structure can continue indefinitely, revealing alternative ways of carving things up: alternative conceptual schemes, logical geographies.

E.g. as more and more is learnt about the architecture of matter, the many forms it can take and the kinds of process that can occur that depend on the structure, we modify and grow our ontologies.

Technology/engineering can contribute substantially to showing what is possible, partly by developing devices that expand what we can observe (e.g. microscopes), and partly by creating things that had never previously existed, that challenge old ontologies.

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0703

EXAMPLE: Types of physical stuff

Concepts of kinds of stuff as those concepts were used before the discoveries about the architecture of matter in the last two centuries.

sand, mud, water, iron, charcoal, wood, air, stone, iron, ...

The new architecture generates a sort of grammar, at different levels:

Possible chemical elements (combining protons, neutrons, electrons, ...)

Possible chemical compounds (more being discovered and created all the time)

links to ontology / mereology

Also new concepts relating to kinds of physical and chemical process (e.g. chemical, sub-atomic changes, etc.)

A grammar specifies what can occur, not what will occur when.

In Chapter 2 of my 1978 book written at Maggie's instigation, The Computer Revolution in Philosophy: Philosophy, Science and Models of Mind, I argued that whereas most people think of discovery of testable, predictive **laws** as the main form of scientific advance (e.g. as argued by Popper), there is a deeper form of advance, that extends our understanding of what is possible by extending our ontology.

Most (maybe all) of the major advances in science have included such changes.
See `http://www.cs.bham.ac.uk/research/projects/cogaff/crp/chap2.html`

Postulates about what is possible are typically not refutable by any experiment, so popular (Popperian and neo-Popperian) conceptions of what makes a theory scientific are mistaken.

A theory about how various kinds of physical structure and process can be rearranged, provides an alternative to "possible worlds" analyses of "can" and "cause".

EXAMPLE: Kinds of "mind-stuff"

Before recent discoveries about architectures of minds (plural), e.g. in the last two centuries, analyses of concepts of kinds of mind, mental state, mental process, mental contents, were restricted by what was known then.

Natural minds vary more than many philosophers have noticed

Humans are not all "normal" or adult: there are also clinical and developmental discoveries

- infant and toddler minds are different from adult minds
- not all adult minds are the same – and some are highly abnormal, e.g. damaged, diseased, etc.

Artificial minds also vary more than most people assume

We have only begun to scratch the surface: a surface Maggie has studied in greater breadth than most people?

The more general space of possible information-processing systems (design-space)

There are many information-processing architectures – natural and artificial

each supporting its own space of possible states, events, processes, contents, ...

which can be divided up in different ways, for different purposes
(every day life, medical, educational, sociological purposes ....)

Each architecture defines a logical topography, supporting alternative logical geographies

A largely unnoticed space: the space of **possible sets of requirements (niche-space)**
We can learn much from engineers: their attempts to design can reveal unobvious requirements.

Don't study only (adult, normal) human minds

Doing philosophy of mind, or psychology, or cognitive science, by studying **only** adult human minds is like trying to do physics by restricting yourself to the surface of the moon.

Even among humans there are many kinds, including

infants, toddlers, teenagers,

people with Down syndrome, Williams syndrome, cerebral palsy, born deaf, born blind, born limbless, or with various kinds of brain damage, or atrophy.

But don't forget microbes, insects, many other invertebrates, etc. etc.

They all have minds in the minimal sense of being capable of selecting among behavioural alternatives on the basis of available information.

They are all **control systems**.

Living things use **informed** control.

This does not presuppose rationality.

We can study living functional control systems using the designer stance (McCarthy) without adopting the "intentional stance" (Dennett), or the closely related "knowledge level" (Newell) except in special cases.

Information processing does not require rationality

From this perspective constantly asking for **rational** explanations of intelligent behaviour, as some thinkers tend to do, would be like a physicist assuming every physical process must involve one object orbiting another.

Most animals are neither rational nor irrational: yet they process information – they acquire, use, store, transform, modify, combine, and in some cases communicate it.

For a discussion of what information is, see:

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

For most organisms, evolution removed the need to be rational, by providing genetically determined mechanisms for selecting among possible alternatives in all circumstances.

Adopting the intentional stance towards their information-processing is a form of anthropomorphism.

Even for those capable of rational decision making, many choices are based on mechanisms either produced by evolution, or by training, or where the mechanisms operate irrespective of whether the choices are rational or not,

For example, if you are going to put your shoes on, it need not matter which you put on first, as long as you start with one or the other: so the selection can be automated.

In some cases choices are irrational, e.g. addictive behaviours, phobias, obsessions, whims, etc.

Exploring designs that **work** leads to a richer logical topography for minds (including more pathologies) than focusing on what can be explained by the intentional stance.

Example: many people assume that motivation must be reward-based but a designer can show how it can be architecture-based.

A simple architecture-based motivational system

Motivation need not be reward-based – it can be architecture-based.

The figure illustrates, in very simplified form, how an architecture-based motivational system might work.

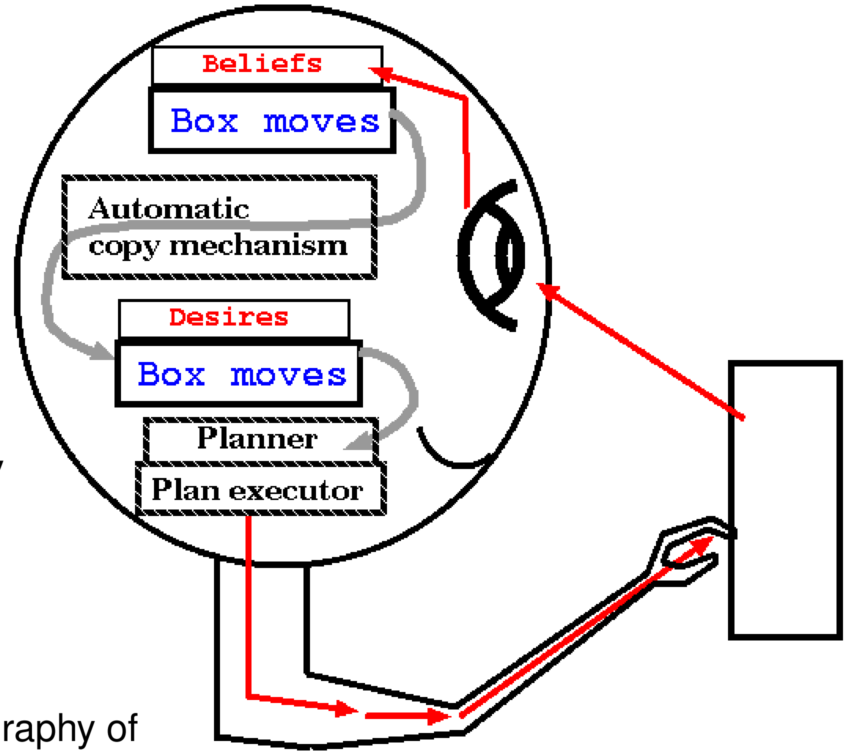

Percepts may trigger formation of new beliefs (descriptions treated as true).

Some of the beliefs (e.g. about processes that have occurred) may trigger generation of related desires, e.g. desires to replicate or to reverse the observed process.

It is also possible for motive generators, motive evaluators/comparators, and other mechanisms concerned with processing of motives to be automatically generated or modified by mechanisms produced either by robot designers or by human evolution – they don't all have to be driven by rewards for the individual.

The diagram is not a serious proposal for a biological architecture: many things are missing.

It merely illustrates the fact that exploring the logical geography of information processing architectures and the states and processes they can support, reveals the need for enrichment of our conceptual scheme: changing the logical geography.

For further (still incomplete) discussion see

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html
Architecture-based motivation vs Reward-based motivation

Metabolism or information processing outside the body?

For most organisms, including humans, metabolism is mostly contained within the body (including the skin), but, for some organisms, part of the metabolism is shared with a host organism or with some other part of the environment.

Examples include animals that begin digestion of food before consuming it (e.g. house flies, and the cooking done by humans), various parasites, symbiotic partners, and birds that make use of air currents – using small amounts of their chemical energy to control their use of much larger amounts of external energy.

A great deal of information-processing can also occur outside the body, and often does.

For example, humans can reason, remember, and refer with the aid of: external diagrams, devices, and other people; and many animals use scents, markings, landmarks, etc. in controlling behaviours.
(But see confusions about embodiment, below.)

There are several different reasons why some information-processing has to be exosomatic, including

- the need to enhance sensors with other devices, e.g. telescopes, microscopes, thermometers, etc.;
- inadequate short term memory for some tasks, e.g. long division, design of complex machines;
- need to use maps for planning routes, etc.
- the reliance on others for identification of spatially or temporally remote individuals whom we have never met (e.g. Julius Caesar (see P.F.Strawson, *Individuals* 1959)).

(See Chapters 6 and 7 of **The Computer Revolution in Philosophy**
`http://www.cs.bham.ac.uk/research/projects/cogaff/crp/`)

NOTE 1: Information does not have to be true

NOTE:

Some philosophers (e.g. Dretske) assume that the label “information” can only be applied to truths, e.g. because if I tell you something false I have not given you information, only misinformation.

Restricting “information” to truths ignores questions, commands, hypotheses, counterfactual reasoning, and of course error.

The universe contains

- Matter
- Energy
- Information (of many kinds, including control information, associate with information-users of many kinds).

Their interrelations are not yet all understood (e.g. in the depths of modern physics).

Note 2: Forms of representation

Some philosophers assume that everything thought, presupposed, conjectured, inferred can be expressed in language, and that all our conceptual mechanisms must operate on something like linguistic structures.

Not all philosophers believe this all the time: e.g. J.L.Austin "Fact is richer than diction" (in "A plea for excuses").

This assumption is hard to square with some of the facts about the intelligent behaviour of pre-verbal humans, and several other species of animal.

I have argued that human language must have precursors used for perception, reasoning, motivation, planning, control of action.

These pre-cursors were used only internally at first.

Jackie Chappell and I call them "generlised languages" (GLs). See

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang`
Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)

I think it is unhelpful – even seriously misleading – to refer to these as involving "non-conceptual content". That implies a bigger difference from e.g. propositional content, than I think exists. E.g. they can include structural variability and compositional semantics. (See the GL presentation.)

I would rather say "non-propositional" or non-Fregean (in the terminology of my 1971 paper).

Something Maggie did to me

She made me take seriously the question whether having a living body was a requirement for having a mind.

- For a long time, I thought having a living body and biological brain, as opposed to having an artificial body and brain was just an implementation detail, that might affect specific contents of mental states and processes, but not whether they existed at all.
- Moreover, I was pretty sure having a body with any particular structure was not a requirement for some of the kinds of thought processes I was familiar with in doing mathematics (e.g. lying on back with my eyes shut thinking about a problem in transfinite set theory).
- Reading what she wrote about autopoiesis, metabolism, e.g. the work of Maturana and Varela, made me start thinking about her criticisms of allegedly intelligent “tin cans”.
- Then one day I had flu, and realised that my mental state could be transformed in all sorts of subtle ways by the activities of nasty little beasties interfering with my metabolism....
- That made me wonder about the extent and role of coupling between low level biochemical processes and high level cognitive functions, and whether perhaps the latter depend in some crucial way on the former, as Maggie had at least hinted they might.
- But what is metabolism?
  I am no expert on this subject. Unlike Maggie, I have just scratched the surface ....

What is metabolism?

Using google I reached the conclusion that the label “metabolism” had different definitions in different research groups.

However it is not hard to find examples of processes in animal bodies that would be described as closely connected with or constituting metabolism according to one or more definitions.

- growth, repair, digestion, distribution of chemicals (e.g. products of digestion, respiration, and waste materials), conversion of chemical into mechanical energy, defence against attacks of various kinds
- These all involve bits of the body being involved in (mostly molecular scale) physical and chemical processes concerned with creating, servicing, or controlling bodies and their functions.

They all seem to involve some form of information-processing, because they all involve control, and control needs to be informed, so as to be sensitive to needs, opportunities, constraints and risks.

The label “self-organising”, which refers to consuming, producing, transforming and rearranging matter and energy, also implies control: some sort of self-control.

The information-processing viewpoint – we study:

- control - selecting among alternatives, according to changing requirements
- This involves using information, in discovering options, selecting between options and in carrying out decisions.
- The simplest case: homeostasis (negative feedback controller)
- Evolution discovered requirements for steadily adding new forms of sophistication.

Beware of dichotomies

Philosophers and others have got used to ontologies (logical geographies) making heavy use of dichotomies: division of everything into two categories, but that can be counter-productive.

In particular, don't look for some binary divide between self-organising and non-self-organising systems: there are many different cases, from water running down a hill to an embryo growing itself.

Likewise do not expect a sharp divide between things that do and things that do not process information.

There are also many different cases, between objects that merely react to forces by accelerating, though negative feedback control systems (homeostats), and other systems that use quantitative information to control quantitatively varying effector signals – to mechanisms that build and use complex enduring information structures with information about themselves and the environment, past, present and future, nearby or remove.

For further subdivisions (not a complete catalogue) see

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604`
Requirements for a Fully Deliberative Architecture (Or component of an architecture

Or things that do and do not have autonomy/free-will etc

`http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#8`
How to Dispose of the Free-Will Issue, AISB Quarterly,1992
(Also expanded as Ch.2 of Stan Franklin's *Artificial Minds*, 1995

Life and Information-Processing

Life takes many forms, But they all involve control — Informed control.

Control involves selection among sets or ranges of possibilities.

Control can be informed by results of various kinds of external and internal sensing.

Initially the control was only of physical processes (online or ballistic).

Later the control processes needed to be controlled – e.g.

- switching on various forms of control at different stages of growth, development, feeding, reproducing, ... partly in response to changing internal needs, partly in response to external changes
- modified by learning –
- evolution can solve only a subset of the control decisions of individual organisms, though it solves surprisingly many, even in humans, when innate (genetically determined) reflexes control behaviour.
- Sometimes evolved or learnt control strategies can go badly wrong, e.g. auto-immune disease, cancers.

Biologists discussing evolution mostly focus on changes in morphology (macro and micro) and behaviour.

In recent years they have increasingly been investigating evolution of information-processing mechanisms, forms of representation, architectures, ...

Information-processing requirements can also change (evolve).

But neither designs nor requirements are in fossil records.

See Chappell and Sloman in IJUC 2007:

Natural and artificial meta-configured altricial information-processing systems

So Life requires Mind

There are debates about whether Mind requires Life.

However, if we construe “mind” as a label for a collection of information processing capabilities – of any kind, then we can conclude

Life requires Mind

Simple life forms have very simple minds – but as needs become more complex so do the control systems.

Arguing about terminology, or simply trying to sort out how our existing systems of concepts divide things into various categories (exploring logical geography) is an activity of limited scope and interest.

We need to broaden philosophical activity to include

analysis of the space of possible sets of requirements (niche-space)

analysis of the space of possible designs for behaving systems (design-space)

investigation of empirical and analytical relationships between possible niches and possible designs

I.e. philosophy needs to move beyond studying logical geographies and do more studying of logical topographies uncovered through scientific and technological advances.

We can then investigate imposing alternative logical geographies on those topographies for different uses

(just as we do for geographical areas on planet Earth).

All organisms are information-processors but the information to be processed has changed and so have the means

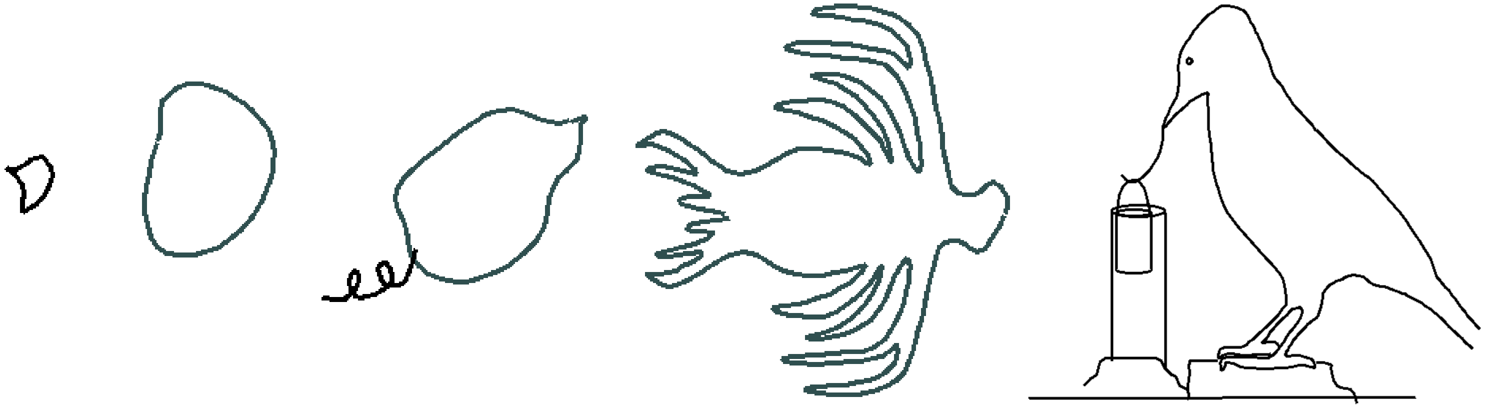

Types of environment with different information-processing requirements

- Chemical soup
- Soup with detectable gradients
- Soup plus some stable structures (places with good stuff, bad stuff, obstacles, supports, shelters)
- Things that have to be manipulated to be eaten (e.g. disassembled)
- Controllable manipulators
- Things that try to eat you
- Food that tries to escape
- Mates with preferences
- Competitors for food and mates
- Collaborators that need, or can supply, information.

The role of the environment

Ulric Neisser:

"We may have been lavishing too much effort on hypothetical models of the mind and **not enough on analyzing the environment** that the mind has been shaped to meet."

Neisser, U. (1976) *Cognition and Reality*, San Francisco: W. H. Freeman.

Compare: **John McCarthy**: "The well-designed child"

"Evolution solved a different problem than that of starting a baby with no a priori assumptions.

....

"Instead of building babies as Cartesian philosophers taking nothing but their sensations for granted, evolution produced babies with innate prejudices that correspond to facts about the world and babies' positions in it. Learning starts from these prejudices. What is the world like, and what are these instinctive prejudices?"

`http://www-formal.stanford.edu/jmc/child.html` Also in AI Journal, December 2008

All biological organisms are solutions to design problems that cannot be specified without specifying in detail the relevant features of the environment.

Turing, surprisingly got this wrong: he thought human-like learning was possible from a "clean slate".

J.J. Gibson understood the general point, but missed many important details.
Other relevant authors: Piaget, Fodor, Chomsky, Mandler, Keil, Gopnik, Tenenbaum, Thomasello, Karmiloff-Smith, Spelke, ...

Confusions about embodiment

There has been much confusion about the relation between information-processing requirements and embodiment.

A common assumption among “embodiment theorists”:

Cognition is driven by (a) body morphology and (b) the sensorimotor interfaces between body and environment, and all learning is concerned with modifying the sensory-motor interaction.

There is another possibility:

Various more or less abstract features of the environment can drive information processing developments, e.g. acquiring and storing information about the layout of the environment on various scales, and developing theories about un-sensed features of objects, detectable only indirectly, e.g. rigidity, edibility, medical properties.

Such information can be acquired by organisms with very different body morphologies, sensorimotor interfaces and even brain mechanisms: e.g. primates, corvids, octopuses. See
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/embodiment-issues.html`

Conjecture: There can be convergent evolution at virtual machine level with very different physical substrates.

Question:
Perhaps even more diversity is possible, including similar minds at virtual machine levels with totally different low level infrastructure?
Maybe even without metabolism?

In that case

Mind does not require life

but

Life does require Mind (information processing competences).

Further investigation should not be based on our current conceptual prejudices (current logical geographies) or even our ethical views.

Instead, with open minds we must investigate what sorts of designs are possible, and which sorts of designs match which sets of requirements (niches)

Leibniz said: let’s calculate.

I say: let’s collaborate

The environment as a driver of intelligence

Biological evolution produced many marvellous products, many of which still cannot be matched by products of human engineering,
e.g. power-weight ratios in motors using chemical energy.

- The solutions were not explicitly communicated from outside to be encoded in genomes, like teachers reciting recipes for children to memorise and follow.
- Likewise the environment can "teach" individual learners to acquire new competences, not by telling them solutions to problems but by confronting them with problems that produce self-modification.
- In young humans, these self-modification processes produce deeper results than any explicit instruction can do.
- One of the new competences that eventually emerges is the ability to learn from instruction: we don't know what has to change in the information-processing architecture to make that possible – the banal answer "the change is learning to use language" does not help.
- You can't learn anything by being told, or by imitation, that you previously lacked the resources to represent.
- The mechanisms that enable self-modification to result from problem-solving are probably at least partly produced by evolution (pre-configured), and partly learnt (meta-configured) (Chappell&Sloman 2007)

Implication for philosophy of mathematics

The role of the environment in posing problems that can drive evolutionary development, including development of cognitive mechanisms and competences was recognised clearly by Neisser(1976) and McCarthy(2008) (see previous slide).

The role of problems posed by the environment in driving individual development, including mathematical development, may not yet have been fully appreciated. (Educational policies can screw this up by trying to impose an unworkable linear ordering on the process.)

There are partial appreciations of these ideas: e.g. Vygotsky's "zone of proximal development", Polya's advice on what sorts of problems to give mathematical learners, and the work of John Holt. I think the work of David Tall on mathematical learning is also relevant, though I have only recently encountered it, and know little about it: See `http://www.warwick.ac.uk/staff/David.Tall/`

For much more on this see

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler`
A New Approach to Philosophy of Mathematics: Design a young explorer, able to discover "toddler theorems" (Or: "The Naive Mathematics Manifesto").

Main theses – aiming to unify a lot of empirical data

- Humans (+perhaps some other species) acquire various kinds of empirical information (e.g. by playing) then LATER develop a new derivation of the information, removing or changing its status to non-empirical – having a kind of necessity (defined later).

  They may not realise they are doing this, unless they study mathematics – perhaps not even then!

- The forms of representation, perceptual and other mechanisms, ontologies, and information-processing architectures that make the transformation possible, evolved to meet biological requirements imposed by complex and changing 3-D environments.

  I'll try to show how this is important for the ability to produce creative solutions to novel problems, without having to do inherently slow statistical learning or dangerous testing where errors can kill.

- The competences required for this are not all present at birth: they develop in layers driven partly by the environment. See Chappell & Sloman

  Natural and artificial meta-configured altricial information-processing systems, IJUC 2007
  `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609`

- In humans, those biological competences provide the basis of the ability to do mathematics, and some of that ability exists (unrecognised) even in young children.

  I'll introduce the notion of a "toddler theorem" and give examples.

- Understanding the biological origins of mathematical competences provides support for Immanuel Kant's philosophy of mathematics,

  wrongly thought to have been refuted by the discovery that physical space is non-Euclidean.

- **I hope the ideas can be tested and demonstrated one day in a "baby" robot.**
- I welcome collaborative research in this area. Any developmental psychologists interested?