Genomes for self-constructing, self-modifying information-processing architectures

Presented at
SGAI 2010 Workshop on
Bio-inspired and Bio-Plausible Cognitive Robotics
http://www.cs.bham.ac.uk/~jlw/cognitive-robotics-workshop.html

Aaron Sloman
School of Computer Science, University of Birmingham
http://www.cs.bham.ac.uk/~axs/

(Some of these ideas come from Margaret Boden, Steve Burbeck, Jackie Chappell, John McCarthy, Marvin Minsky, and no doubt many others.)

(Work in progress – liable to be updated.)

The latest version of this presentation is in my 'talks' directory:
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk89

Abstract

Between the earliest proto-organisms and modern complex animals, evolution has produced hugely varied physical forms, physical behaviours, life-cycles, and modes of reproduction. All use information of various kinds (along with matter and energy). Information contents control internal and external processes: including changes through learning and development of individuals, and in genome development by evolution. Information contents are also involved in social and cultural changes.

I conjecture that animal intelligence involves far more diversity in forms of information-processing than we have dreamed of so far, and that major sources of that diversity lie in various layers of complexity in the physical, biological, and social environments in which evolution and individual development occur. I also conjecture, contrary to assumptions of many researchers interested in embodiment, that some of the more abstract features of the environment (e.g. rigidity of some materials and diversity of 3-D terrain structures) can have a common influence on evolution of information-processing in animals with very different neural mechanisms, sensory and motor systems and morphology, though implementations of common functionality can differ enormously. (This is an unnoticed implication of some of the points in McCarthy (2008).)

The ontologies currently used for observation and theory-construction by most neuroscientists, psychologists, biologists, and AI/Robotics researchers cannot (as far as I know) accommodate all of the "exo-somatic" aspects of animal ontologies, seriously restricting the explanatory power of theories using those ontologies. These requirements also do not fit some of the most common ways of thinking about "embodied" cognition (Sloman, 2009c).

I shall address some ways of overcoming those limitations, in part by illustrating environmental features that have received little attention, and in part by decomposing some of the functions of a genome.

Background Information

The talk builds on ideas in these two papers (among others):

Jackie Chappell and Aaron Sloman,
Natural and artificial meta-configured altricial information-processing systems, in
Int. J. of Unconventional Computing. vol 3, No 3, 2007 pp. 211–239,
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609`

Aaron Sloman,
What's information, for an organism or intelligent machine? How can a machine or organism mean?,
In *Information and Computation*, Eds. G. Dodig-Crnkovic and M. Burgin, World Scientific, 2010,
`http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#905`

And presentations here:

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#wonac07`
Causal competences in animals and machines (With Jackie Chappell).
(Including Humean and Kantian causal understanding.)

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk35`
A (Possibly?) New Theory of Vision

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang`
Talk 52: Evolution of minds and languages.

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#brown`
Talk 68: Ontologies for baby animals and robots

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cafe04`
Talk 28: Do Machines, Natural Or Artificial, Really Need Emotions?

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk84`
Talk 84: Helping Darwin: How to Think About Evolution of Consciousness

And others in the same web site.

Maynard Smith on Information in Biology

After preparing most of these slides I stumbled across and read this paper by John Maynard Smith (JMS):

The Concept of Information in Biology (Maynard Smith, 2000)

That paper is broadly compatible with the ideas presented below: but JMS was mostly concerned with genetic information, whereas I am trying to draw attention to a much wider range of roles for information in organisms, many of which are also relevant to robots, for example, the roles of information contents in perception, learning, problem-solving, motivation, planning, and generation and control of actions.

JMS focuses mainly on what he calls "symbolic" encodings (partly inspired by Monod): "...it is the symbolic nature of molecular biology that makes possible an indefinitely large number of biological forms", denying that there can be necessary connection between a form of expression and the information content.

He also follows Dennett (1987) and others in claiming that information must be intentional in the sense that an encoding of some information exists because it was created or selected for the purpose of communicating that information, and the selection is arbitrary because there is no intrinsic relationship between encoding and content.

These criteria for use of information are too restrictive, for reasons previously explained in Sloman (2008a). For instance, they fail to allow for the information that a perceiver can get from the environment, which need not be intentional in that sense, since many objects that provide information about themselves were not created or evolved for that purposes. Encodings are not always arbitrary, e.g. there are deep connections (but not isomorphism) between the information content a seen object provides about itself (by reflecting light) and the structure of the object. Maps are useful because of their structural relationships with what they represent (though, as pointed out in Sloman (1971), those structural relationships can take many forms). In a mathematical or logical proof, there are also necessary structural relationships between the formal expression of the proof, and the content of the proof, even if some of the notational details are arbitrary.

A point in JMS's paper that is central to what follows is this: "The picture that is emerging is one of a complex hierarchy of genes regulating the activity of other genes". We'll see that the processes are more varied and more context-sensitive (or parametrised) than that wording suggests. (Compare Hawkins (2011))

How to understand a multi-layered genome

- It is very hard to understand a very complex mechanism if you don't have a good theory about what it is for, in addition to observational evidence of what it does and does not do in various situations.

  (Observation does not necessarily provide information about most of what it **can** and **cannot** do.)
- It's even harder for mechanisms built up in layers of functionality, over millenia, with new layers of functionality built on old layers – including rarely required but important abilities (e.g. for mating).
- Trying to find out what a mechanism does, or how it works, without having a good theory to guide both observation and theory construction can lead to serious gaps and errors, in observations and conclusions.

  (Like future archaeologists discovering an old mechanical clock with half-day spring, and deciding that its main function is to produce a regular soothing ticking sound, and the role of the clockface and hands is to indicate when it will next need rewinding, with an ingenious mechanism for indicating the gap with two levels of precision.)
- The task is more difficult if some mechanisms (e.g. vision in animals) serve multiple purposes, sometimes simultaneously sometimes in sequence (Sloman, 1993).
- One way to tackle the complexity of the scientific task of understanding a particular highly complex genome is to try to identify (implicit) "design decisions" taken by evolution on the way to the present genome, and also to compare it with a wide variety of other biological designs, supporting partly different functions.
- In particular, I'll focus on the information-processing design decisions, since others have already received much attention. (NB: Designs don't need intentional designers.)

Biological control by information

The earliest and most basic biological use of information is for **control**.

The boundary between physical control and control by information is fuzzy in simple cases (e.g. simple homeostasis) – but becomes clearer as differences between information content and physical causes increase: e.g. using information about the past, about remote unperceived objects, or about non-existent, but possible, future states.

Examples of **physical** causation of animal behaviour:

- an elephant unwillingly slithers down a muddy slope – pulled by gravity.
- an ape falls from a branch after losing its grip and on the way down knocks some fruit to the ground.

Examples of **information content** playing a role in causing actions:

- an elephant follows a route it remembers.
  Remote portions of terrain in the route are not in contact with the elephant and cannot physically influence the actions. (See `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk86`)
- an ape sees fruit on a branch lower down and uses information about the possibility of knocking the fruit to the ground in deciding to let go of the branch it is holding.

A difference between using energy or matter and using information

- As energy is used, it is consumed, or made unusable. Often (but not always) matter is also consumed as it is used, e.g. using materials to make a shelter, or eating food.
- Much information (e.g. in a recipe) can be re-used indefinitely, without being altered.

NOTE: Piaget noticed the importance of information about possibility and necessity: (1981, 1983)
(Unfortunately some of his experiments presented children with over-complex problems.)

Evolutionary transitions

Between the earliest proto-organisms and modern complex animals, evolution has produced hugely varied physical forms, physical behaviours, life-cycles, and modes of reproduction. All use:

- **matter and energy**
- **information of various kinds**

- Information about what?
- How do they get it?
- What do they do with it?
- How do they put it to use?
- How is the information represented/encoded and manipulated?

There are many different answers to be investigated.

One of the major features of biology is **diversity.**

Understanding evolutionary and other trajectories can help us explain current designs:

studying evolutionary, epigenetic and social-cultural transitions producing all that diversity can help us understand what difference having or not having some feature can make in various circumstances.

DO NOT ASK: What's the right, or best, design?

Many AI/Robotics researchers attempt to understand and replicate human intelligence without taking account of either the diversity of types of human intelligence (e.g. infants, toddlers, etc.) or the many similarities and differences between human and non-human types of biological intelligence: we need to study a wider space of possible designs.

Identifying important transitions in evolution

John Maynard Smith and Eörs Szathmáry (1995) **(JMS&ES)** proposed that there are eight major transitions in evolution, summarised here:
http://en.wikipedia.org/wiki/The_Major_Transitions_in_Evolution

Transition from:	Transition to:
1 Replicating molecules	"Populations" of molecules in compartments
2 Independent replicators (probably RNA)	Chromosomes
3 RNA as both genes and enzymes	DNA as genes; proteins as enzymes
4 Prokaryotes	Eukaryotes
5 Asexual clones	Sexual populations
6 Protists	Multicellular organisms - animals, plants, fungi
7 Solitary individuals	Colonies with non-reproductive castes
8 Primate societies	Human societies with language, enabling memes

They identified features common to these transitions:

1. Smaller entities come together to form larger entities.
2. Smaller entities become differentiated as part of a larger entity.
3. Smaller entities become unable to replicate in the absence of the larger entity.
4. Smaller entities become able to disrupt the development of a larger entity.
5. **New ways arise of transmitting information.**

NB: That view of information-processing is too narrow.

Transmitting is not the only thing done with information.

Why transmit it, if not to enable something else to be done with it? (Examples below.)

Important, under-studied, transitions

JMS&ES mentioned **new ways of transmitting information,** but transmitting is only one of many things that can be done with information, and the value of transmission depends on the uses of what's transmitted.

In many cases, transmission of information would not have had any biological use if the information had not had a prior use.

By identifying other processes, besides transmission or communication, we can begin looking for a wider variety of evolutionary transitions involving information, including:

- new kinds of information content (examples are given later);
- new forms of representation (Sloman, 1971; Peterson, 1996; Minsky, 1992);
- new ways of acquiring information, including
  - new kinds of sensing/perceiving,
  - new ways of using evidence,
  - new kinds of reasoning, inferring, hypothesizing;
- new ways of storing and retrieving information;
- new ways of processing, or manipulating information (parsing, combining, inferring, ...);
- new uses for information (e.g. controlling, predicting, explaining, asking, forming goals, plans,...).

Changes of the above kinds can occur:

- in evolution (across species);
- in development and learning (over time, in an individual);
- in various social interactions (collaborative or competitive) in groups, or societies.

This is a shallow, illustrative, high-level list of change categories! We need details.

Informational life-cycles

Control by information occurs, at different levels of abstraction, e.g. in:

- evolution
  - selection of genomes and genome reproducers, by selecting between their effects;
- reproduction of micro-organisms;
- development of an individual from a fertilised egg;

  including growth of individual information-processing mechanisms (using chemical information) controlled by both internal and external information;
- production of behaviours of individuals and pairs or groups of individuals;
- formation and propagation of cultures.

Examples of information-based control in an individual organism include

- metabolism (– more than just physics and chemistry);
- detecting changed, more global, internal states (e.g. need for food, need to breathe);
- altered perceptual contents (very many sub-categories – see below);
- wholly and partly genetically driven development (compare Jablonka and Lamb (2005));
- individual learning (there are probably more sub-types than anyone has analysed);
- social and cultural processes

  producing effects directly and indirectly, through explicit and implicit forms of communication. Karl Popper's "Third world" includes these (Popper, 1972). "Third world" entities are produced not only by human thought and action, but also by biological processes. (Compare (Hall, 2006).)

These processes use diverse types of information content and forms of representation.

Transitions in **information content**: some samples

Some transitions in evolution and development involve changes in information contents processed – what the information is about – such as transitions:

- from recognition of structures by structures to sensing of scalar values;

  In very simple organisms detection is done by chemical structures interacting – there is no nervous system and probably no measurement of scalar values, which would require more complex sensors. Compare the recognition mechanisms used by immune systems, and cell repair systems.
- from somatic information (about internal states, including contents of sensor and motor signals) to exo-somatic information contents (about things **outside** the organism);
- from sets of measures and labels to structural descriptions (e.g. parse trees Minsky (1963));
- from information about relations between simple objects (e.g. contact) to information about objects that are themselves bundles of relations (e.g. a hand, or plant, or animal);
- from object relations to multi-strand relations between objects with parts (Sloman, 2008b);
- from simple processes to multi-strand processes (many relations changing together);
- from representing only physical entities to representing things that use information –

  i.e. from semantic competences to meta-semantic competences (e.g. representing prey or mates);
- from process information used transiently, e.g. for servo-control, to persistent information about past processes, usable for multiple purposes at later times;
- from representing only transient states and processes to recording enduring spatial structures, not all constantly visible;
- from encoding simple changes to representing spatio-temporally extended processes;
- from particular situations to ontologies, generalisations and theories. (Sloman, 1978, Ch.2)

Fully and partly specified information content

There are many differences between ways in which information can influence what happens, including:

- just causing itself to be copied from one location to another;
- causing information content to be copied, while the form of representation is altered

  e.g. going from a map to a description or vice versa, converting a description or algorithm from one computer language to another, and translating information about what survives into DNA.
- using a specification of some generic type of object or process to produce an instance of that type,

  e.g. assembling a new instance of a type of molecule, or producing a new instance of cell-repair, a new instance of a swimming motion, or a new instance of the elephant genome.

Very often a generic information item **I1** used in the production of a new object or process **does not suffice** to specify all the details of the instance produced.

- Instead **I1** has to be combined with other information items of different sorts **I2**, **I3**, ... to produce the result (e.g. details of size, behaviour,...).
- In some cases, **I1** is parametrised, i.e. has explicit gaps to be filled when it is used.
- In other cases the mechanisms using **I1** are influenced by other items of information:

  The process of using **I1** is then context sensitive, i.e. it depends on other items of information that may be available throughout the process, or become available from various sources during the process of using **I1**. In some cases **I1** itself specifies where to get the additional information.

Much of the control by a genome is **parametrised** or **context sensitive**, including the growth of an organism's control systems, as in JMS's genes regulating other genes.

(Also: controlling language development that is influenced by linguistic and non-linguistic environments.)

Evolutionary transitions in mechanisms of evolution

A more complete presentation would need to show how new mechanisms of evolution, including new forms of evolutionary information processing, can themselves be products of evolution – but what kind of evolution?

- Some evolutionary mechanisms allow new organisms to be slight variants of previous organisms because of (usually small) random changes in genetic materials, as in "The Neo-Darwinian Modern Synthesis" (NDMS).

  Compare mechanisms of individual learning that gradually shape new behaviours by rewarding or punishing small, random, changes in behaviour.

- Endosymbiosis allows substantial pre-formed products of evolution to be combined in new ways, producing discontinuous changes in products of evolution,

  (as summarised in `http://en.wikipedia.org/wiki/Lynn_Margulis`, and also discussed in Hawkins (2011)).

  Compare the ways in which new phrases and sentences can be formed by combining previously constructed words, phrases and sentences, rather than using only small gradual modifications of previous sentences, which would take far longer to generate the variety of utterances produced by any human, with much nonsense intervening between useful thoughts and communications.

- Those mechanisms can be investigated as examples of physical and chemical processes: but full scientific understanding requires them to be analysed as mechanisms for transforming information – which in the case of endosymbiosis requires transformations that recombine complex information structures in new ways.

- Such recombination seems to have been controlled by information-processing mechanisms, long before humans invented selective cross-breeding.

Slow gradual transitions vs discontinuities

The previous slide introduced an important distinction between

- forms of representation that allow only continuous or very gradual change

  e.g. representations using either real numbers or continuously varying physical properties, like voltage, or size, to carry information

- forms of representation that allow large or small changes but where most large changes will produce something unusable

  E.g. using a bit pattern with mechanisms for encoding to and decoding from the bit patterns (as in "genetic algorithms": GAs).

  Arbitrary changes in such patterns will mostly produce unusable junk, especially if changes are large.

- forms of representation that allow sudden changes of arbitrary complexity, where many such changes will produce usable new information contents,
  - E.g. using a grammatical formalism where allowed changes include only well-formed sub-structures being replaced by other well-formed substructures that have already been found to be useful
  - Examples: replacing a referring phrase or process description in a sentence with another possibly more complex one, or replacing a sub-graph in a graph with another sub-graph that has been found useful (e.g. in "genetic programming": GP)

These distinctions between types of change of information structure and information content are relevant to biological changes

- in evolution, across generations of individual organisms
- in learning and development, within an individual organism, producing long term changes in information content
- in momentary and transient processes of perceiving, thinking, asking, planning, deciding,

Piaget seemed to understand much of this but expressed it unclearly: Piaget and *others* (1983).

Transitions involving communication

Most examples of information-processing above have involved uses of information that don't require communication, e.g. perception, learning, planning.

However, many organisms communicate information to other organisms, sometimes passively (e.g. by leaving footprints, or making sounds that give their presence away), sometimes actively: and many evolutionary, developmental and cultural changes involve new forms of transmission, or new kinds of content being transmitted or new uses for transmitted information (e.g. collaborative construction of shelters, or hunting).

Because so much research and speculation has been concerned with evolution of forms of linguistic and non-linguistic communication, I'll deliberately say little about this here

(except to point out that what is often described as a process of language learning should be reconstrued as a process of collaborative problem-solving where the problem is how to communicate, using a new rich shared language, as clearly happened with deaf children in Nicaragua. (Sloman, 2008a))

Understanding the mechanisms and forms of communication needs to be based on understanding the motives, perceptual and behavioural competences, and forms of learning that have to precede any form of linguistic communication (so that there is something to communicate, and so that communication can be useful Sloman (1979)).

Unfortunately many discussions of language learning or language evolution totally ignore the non-linguistic cognitive (e.g. representational) prerequisites for use of language.

I'll return now to **transitions in non-communicative uses of information.**

Why change blindness must have evolved late!

There are very famous experiments demonstrating change-blindness of various kinds in humans.

Many examples, including videos, pairs of pictures with differences, and papers, lectures, discussions, can be found by giving ”change blindness” to google.

- It is often assumed that if a change occurs in a perceived situation or object, then that change will be detected, unless something anomalous prevents detection.
- The opposite assumption should be made: change detection needs to be explained, not change blindness – then we can discuss failures in detection mechanisms.
- Sensing/perceiving mechanisms cannot perform change detection unless they have a richer architecture than is required for simply perceiving, and possibly reacting to changing information.
  - A ball rolling down a helter-skelter is subject to forces whose directions change, causing the ball to be constantly changing its direction of motion.
  - The ball has and reacts to information that is constantly changing, and it reacts by changing its behaviour.
  - But that does not mean that it has or can use the information that forces acting on it have changed, or that its direction of motion has changed.
- Organisms can react to changing information by changing their behaviour, without acquiring or using any information about the occurrence, or direction, or amount of change: they lack the required information-processing architecture, with mechanisms that record past states and compare present and past states.

Ontologies needed for information about change

Evolving abilities to represent and detect changes can have many benefits, depending on type of change and how the information is used.

A very primitive system may react only to particular sensed states: e.g.

- if a noxious chemical is detected, move away from it: a change in what is sensed causes a change in behaviour.
- if light is sensed as coming from a certain direction move in that direction.

Detecting change as opposed to merely reacting differently requires an architecture that retains information about past states, while obtaining information about the current state.

- A new internal sensor could then detect changes in sensed states, by comparing the present state with the previous state (e.g. using a two-place queue).
- The simplest version would merely record that a change has occurred, perhaps triggering a switch to some new defensive or exploratory behaviour.
- A more complex version would record information about types of change
  - X has stopped
  - Y has started
  - State or feature X is replaced by state or feature Y
  - Direction or value of some change in intensity; or rate of change; or acceleration (rate of rate...)
  - if several successive states are recorded, then a **process descriptor** could be constructed

Evolution produced many more sophisticated sensors, ontologies and architectures able to detect and use information about differences, changes and extended processes.

Some organisms can also invent new types of change (e.g. making/using new tools).

Further transitions in change detection

Examples of change detection that require specific new mechanisms to develop, include becoming able to detect:

- a proximal change in external stimulation, e.g. temperature, pressure, presence of some chemical, illumination, etc.;
- a change in the environment beyond the organism, e.g. some object moving, perceived surfaces changing curvature or orientation, wind blowing, darkness falling;
- correlations between two or more changes,
  (e.g. moving in a certain direction is associated with increasing or decreasing some sensed value);
- relationships between transient actions and enduring changes in the environment
  (e.g. pushing a rock moves it to a new location);
- relations between external behaviours (e.g. ingesting food) and internal sensed states
  (e.g. level of stored energy increases);
- altered spatial relations between the organism and fixed structures in the environment
  (e.g. the organism changes its location in the terrain, temporarily);
- changes in previously static extended spatial structures
  (e.g. a path blocked, a new gap appearing in a barrier, a gap getting wider or narrower, a source of food or danger changing its location);
- changes in position, orientation, direction of movement, ... of objects characterised in terms of affordances, such as food, or things to be avoided.

(Compare confused debates about ventral and dorsal visual streams, and "mirror" neurons.)

Many special cases have been studied in humans, other animals and robots.

But they are only samples of a huge, under-explored, space of requirements and designs.

Investigating biological information

A possible large scale collaborative project could aim to produce a comprehensive analytical overview of

- biological information contents and how they differ between species
  – either because of differences in environments, or because of differences in organisms;
- ways of acquiring information and how they differ;
- forms of representation and how they differ;
  – physical forms (chemical, neural, physical structures, ...)
  – virtual forms (implemented in physical forms;)
- types of analysis and interpretation of information in organisms, and how they differ;
- ways in which new information is derived or constructed, and how they differ;
- ways in which information is used – in external behaviour and internal processes, and how they differ.

Related questions arising for each case would include:

- what the information in the genome needs to be, in order to enable individuals to develop such information-processing abilities

  (including information directing construction of genome-interpreting mechanisms!)
- how that genomic information is represented (including interacting sub-genomes)
- how it influences the developmental process (how developmental mechanisms use it)
- whether and how that influence is context-sensitive, i.e. dependent on
  diet, physical demands, play opportunities available, plentiful or scarce food, presence of dangers of various kinds, materials available to interact with, language spoken, etc..

Why I do not attempt to define "information"

This presentation (like many others) constantly makes use of the concept of "information".

But there is no point attempting to offer an explicit definition of it:

Attempts at explicit definitions end up going round in circles via ("meaning", "reference", "semantic content", "significance", "connotation", "interpretation" and others).

Yet the concept is indispensable, and has been for centuries

e.g. when people talked, thought and wrote about news, rumours, reports, stories, history, signals, messengers, information, misinformation, etc.

(This concept of information preceded, and is quite different from, Shannon's syntactic concept of something measurable, related to capacity of communication channels.)

- Although "Information", in the sense we need here, cannot be explicitly defined without circularity, we can still make good use of it (like the concepts "matter" and "energy")
  (For more a more detailed discussion see Sloman (2011b), with criticisms of common definitions.)
- We can develop increasingly complex and powerful explanatory and predictive theories using and implicitly defining "information", and show that this is a progressive, not a degenerating, research programme.
  (In the sense of Lakatos (1980), as explained in Sloman (2007b).)

What is sometimes forgotten is that besides different contents, information items can have different **uses**, including recording, informing, misleading, formulating conjectures, questions, motives, imperatives, recipes, procedures, strategies, etc. All of these (and others) are important in biology.

We need to understand a wide variety of designs

Apologies to people who have previously heard me going on about the need to study "The space of possible minds". Some take the need seriously (e.g. Murray Shanahan (2010)) but not most researchers.

Understanding a design feature of a complex system requires knowing what difference it makes – i.e. comparing systems with and without it.

A research community cannot understand one complex system (or even a small number of systems) without studying differences from other systems:

E.g. building too few types of robots leaves us ignorant about what we have and have not achieved.

So understanding causes of pathological behaviour requires understanding "normal" functioning.

Trying to build a machine to pass the so-called Turing test is of little scientific or philosophical value – it's too specific. Turing himself proposed it not as a test, but a prediction to be defended against objections. (Turing, 1950)). See also "The Mythical Turing Test":
`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-test.html`

Studying "neighbourhoods" in design space and requirements space (niche space) and relationships between both sets of neighbourhoods can help increase understanding. More varied neighbourhoods teach us more. (Sloman, 2000)

E.g. what are the different sorts of requirements for organisms of different sorts (worms, fish, lizards, elephants, birds, climbing mammals) to be able to represent spatial organisation?

E.g. how do requirements for forms of representation of spatial organisation differ between:

- making a clock,
- thinking about making a clock?
- repairing a clock made by someone else?

Relations between spaces and trajectories in them

Possible designs
for working systems

Related to

Possible sets of requirements
(niches)

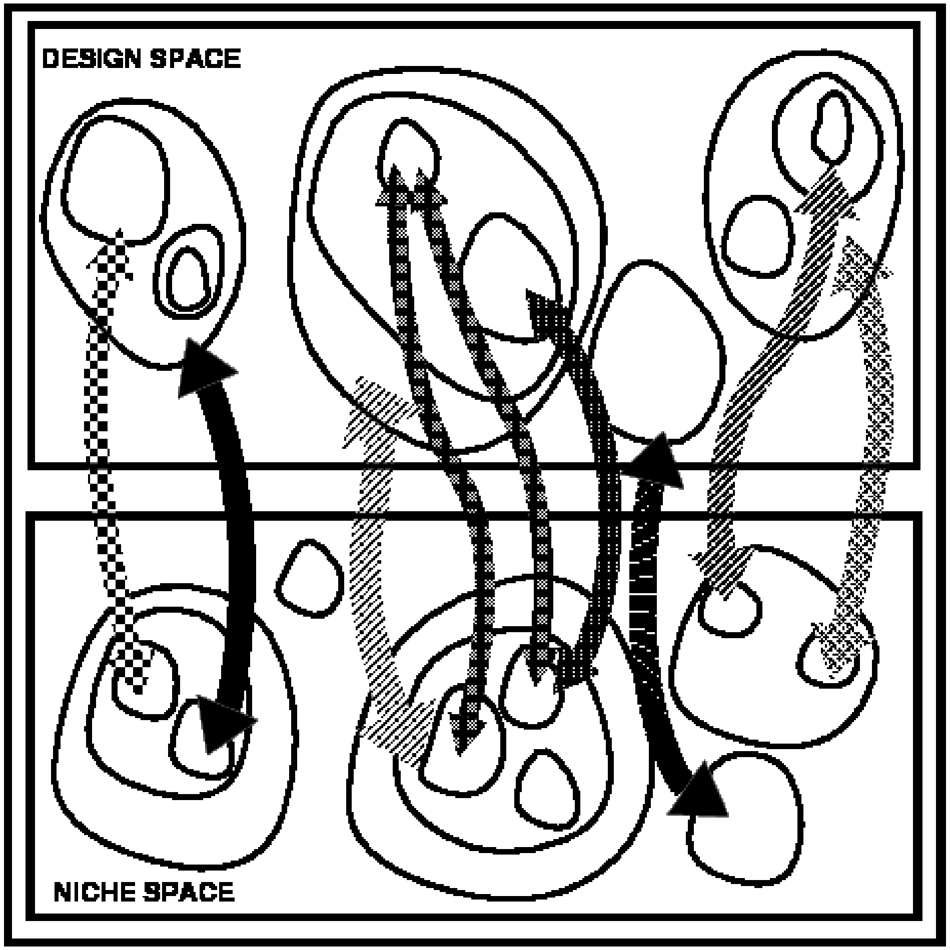

There are many different relations: don't just look for fitness functions.

Compare: the **descriptive** content of consumer reports on products available.

Interacting trajectories occur when there are two-way causal connections between changes in requirements (niches) and changes in designs.

Niches and designs change: There are related trajectories in both spaces

- i-trajectories: individuals develop and learn
- e-trajectories: species evolve across generations
- r-trajectories: (non-biological) a 'repairer' takes things apart and alters them
- s-trajectories: societies and cultures develop (Not shown)
- c-trajectories:
  e-trajectories where cognitive mechanisms and processes in individuals influence trajectories, as in mate selection, or adults choosing which offspring to foster in times of shortage, or selective breeding by farmers.

We shall almost certainly need new kinds of (largely non-numerical) mathematics to model these processes.

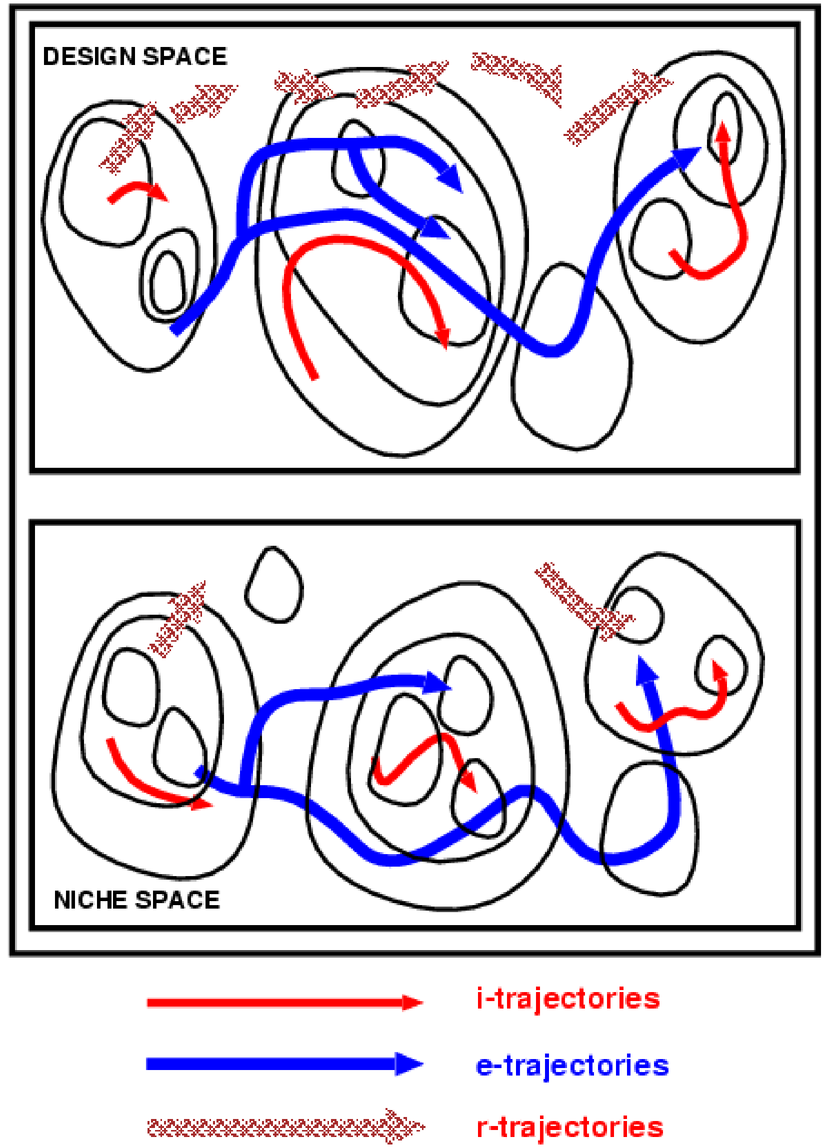

The importance of the environment

Environments must be studied closely, not just the occupants:

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 (1976))

A view echoed in McCarthy (2008)

"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?"

Avoid assuming there is only one environment, in which organisms reproduce, develop and behave, and only one interesting kind of mind produced by evolution.
Corvid minds, primate minds and aquatic mammal minds differ in requirements.

- This planet has provided very different environments at different times and different locations.
- Very different organisms have evolved even in the same environment.
- Not all species have the same kind of relationships between genome and competence: most animals are precocial: born or hatched with all the perceptual and behavioural competences they will ever need (except for minor adaptive changes in parameters – Piaget's "assimilation"?).
- But some species, e.g. humans, are altricial: developing many of their main competences after birth. (Piaget's "accommodation" and more. Some re-build their "epigenetic landscape" as they develop.)
- Why don't those follow the pattern of precocial species and start with all their main competences – though they start with some competences pre-configured, others meta-configured?
  See: Sloman and Chappell (2005), Chappell and Sloman (2007).

Transitions in information-processing requirements

Changing relationships between organism and environment

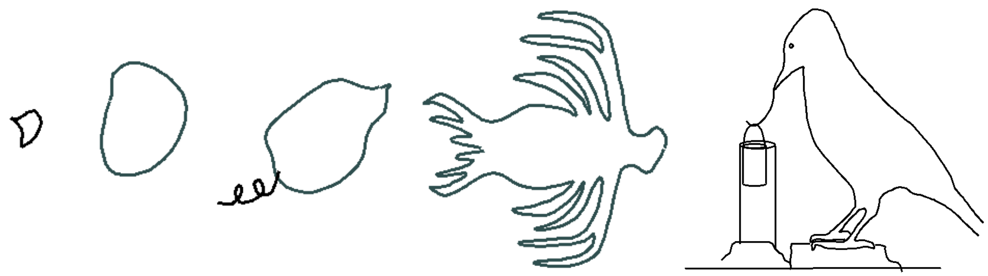

Types of environment with different information-processing requirements

- Microbes in chemical soup, can be wholly dependent on nutrients in their neighbourhood.
- Soup with detectable gradients: offers opportunities to improve location if motion can be controlled.
- Soup plus some stable structures, e.g. places with good stuff, bad stuff, obstacles, supports, shelters, offers advantages for organisms that can build up long term spatial memories, and can plan routes.
- Things that have to be manipulated to be eaten (e.g. disassembled) need new forms of process perception and process control, including ontologies that include kinds of “stuff” with different properties.
- Controllable manipulators: mouth, hands, feet, require different uses of information.
- Things that try to eat you – lead to forms of control for escaping, hiding, defending, etc.
- Food that tries to escape – lead to forms of control for chasing, trapping, lying in wait, heading off, etc.
- Mates with preferences, lead to forms of control of behaviours for attracting attention and winning favour.
- Competitors for food and mates lead to deception, fighting, warding off, defending territory, etc.
- Collaboration with others requires action controlled to aid collaboration, and possibly communication.
- and so on (including information-processing in plants) .....

How information-processing **requirements** change across such cases, depends on both features of the environment and features of the organism (products of previous evolution). Contrast: how **designs** change.

Organisms do not start off ready made

They build themselves

Each is alive and functioning from the start, till death – unlike machines –
(though some need external life-support systems temporarily, e.g. parent.)

They don't just build their physical bodies.

They also construct their own information-processing systems:

- all the way up from molecules,
- collections of connected co-operative sub-structures
- with various neural and chemical (molecular) information-transfer systems, control systems, and information-storage systems,
  (some of them multi-functional subsystems.)
- operating to a large extent in parallel.

Some of the changes seem to require new physical designs
(e.g. late development of specialised brain sub-structures),

other developments depend on changes in use of flexible existing subsystems
(e.g. learning a new motor skill or new branch of mathematics?)

Humans (perhaps not all – but many) go on extending their information processing capabilities throughout life.

Some of these are architectural changes, often in virtual machinery:

– new subsystems are added

– new connections are made

E.g. new kinds of self-monitoring.

Examples of genome-influences in development

Directly or indirectly, genetic information controls all of the following,

though much of the control is shared with other influences

(hence the control is **parametrised**, **context-sensitive**, or **"regulated"**):

- copying DNA
- producing various types of RNA
- synthesising proteins
- assembling cell structures
- driving and modulating the above processes as well as processes in complete cells (metabolism, digestion, damage detection and repair, excretion, growth)
- controlling development of external and internal morphology
- controlling growth of sensory and motor subsystems, and their connections
- controlling initial forms of behaviour (e.g. breathing, sucking, initial movements)
- generating and controlling experimental/exploratory behaviours (Rochat, 2001; Gibson & Pick, 2000)
- controlling post-natal changes:
  e.g. acquiring specific information, behavioural competences, new ontologies, new generalisations, new theories about how the environment works Chappell and Sloman (2007)
- controlling production of motives/preferences/values
  (under multiple influences, in evolution, learning history, recent sensing, etc. (Sloman, 2009a))
- enabling more and more control processes to be shared with the environment.
  E.g. using compliance of supports, getting help from gravity, from currents, from up-drafts.
- There are enormous variations
  – across species, subsystems, stages of development, and across individuals at similar stages.

The more abstract the influence the more obscure the relationship with physical structure and processes

Human scientists and engineers have discovered the benefits of using more and more powerful virtual machines whose powers and properties are increasingly remote from those of the underlying physical machinery.

- There are several reasons to suspect that similar benefits from use of virtual machinery were "discovered" much earlier by biological evolution.

  `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk84`

- In that case it may be practically impossible to discover what the machinery is –

  either bottom-up, by studying the chemical and neurophysical details in the hope of finding their functional powers, **or outward-in** by seeking correlations between physically detectable internal states and processes and externally observable behaviours in different contexts.

- If so, such bottom-up and outward-in research strategies need to be combined with creative (informed) searching (top-down) in abstract design spaces.
- That cannot work unless we have learnt enough about powerful building blocks for virtual machinery. (We need powerful ontologies and theories.)
- That can still leave impossibly complex search spaces: but we can constrain searches by investigating requirements that the designs have to satisfy.
- That can be facilitated by doing comparative analytical studies of capabilities, across species, across individuals and across time, or between contexts within an individual.

  Contrast researchers seeking only **commonalities** and **correlations** between observable behaviours and situations, without theories about information-processing mechanisms.

The methodological significance of transitions

The notion of "transition", applicable in several different contexts and at different levels of abstraction, can help us understand the varying relationships between genomes and information-processing architectures.

Maynard Smith and Szathmáry discussed eight types of transitions in evolution. But there are many more relevant to the evolution of information processing capabilities (minds):

- transitions in the niches, or sets of requirements that play a role in natural selection, via changes in the environments in which organisms live, compete and die
  (including changing information available/unavailable about constraints and opportunities);
- transitions in the information-processing competences and designs that are relevant to coping better or worse with those requirements;
- transitions in the implementations of those designs, since typically any abstract design can be implemented in many different ways, and changes in implementation may not affect the original requirements, but may alter meta-requirements,
  e.g. robustness or possibilities for deeper integration, or for future changes (Sloman, 2007a);
- developmental transitions in individuals – changes in requirements, designs, and implementations during development – supported indirectly by the genome.
  (Requirements for infants, toddlers, teenagers, parents, etc. all differ.)

There are tradeoffs between processes of evolution and processes of individual development.
The transitions in both may be closely related, as we'll see.
Pressure for them comes partly from the environment, partly from previously evolved forms of embodiment.

Studying interactions between transitions, with increasingly complex feedback loops, may explain some drivers of evolution and development, leading to a better understanding of the products of both.

Meta-semantic competences – who has them?

Semantic competence

Being able to represent, reason about, perceive, act in relation to things of various kinds, using information about them.

Meta-Semantic competence: (a special case of semantic competence)

Being able to represent, reason about, perceive, act in relation to, things with semantic competences: i.e. being able to use information about information-users.

Is it obvious which organisms have meta-semantic competences?

(Compiled innate infant reflexes may give the misleading impression of meta-semantic competence.)

A special case of meta-semantic competence is Meta-cognition:

Being able to represent some of one's own semantic competences and their applications – e.g. detecting, reasoning about and manipulating:

- sensory contents
- intentions, goals, preferences
- beliefs and memories of various kinds (individual and general)
- procedures used in deciding, reasoning, planning, etc.

Issues:

- Self-related and other-related meta-semantic competences
- Complexities of meta-semantic competence
  - How to deal with referential opacity.
  - Use new logics or use new architectures? The latter!
  - Can brain-imaging methods give clues? Probably only after we have good theories.

See Sloman (2011a).

Informational biodiversity and commonality
(convergent evolution)

I conjecture that

- animal intelligence involves far more diversity in forms of information-processing than we have understood so far, and that
- major sources of that diversity lie in various layers of complexity in the physical, biological, and social environments in which evolution and individual development occur (many of them still undocumented).
  (There is also much to be said about information-processing in plants.)

Despite the diversity, there are commonalities (convergent evolution):

- some of the more abstract features of the environment (e.g. rigidity and diversity of 3-D terrain structures) can have a common influence on evolution of information-processing in animals with very different neural mechanisms, and very different sensory and motor systems and morphology. (E.g. birds and mammals.)
- however implementations of common functionality can differ enormously, illustrated by detailed differences in several species that are able to perceive, manipulate, and move in relation to 3-D spatial structures on various scales:
  - primates
  - birds
  - cephalopods (octopus intelligence is full of surprises.)
- For deep understanding we shall have to learn both about the common functions and the different implementations: each will help to constrain the search for the others.

An example – ontologies for gripper-users

Consider an animal with grippers that can move independently of each other and independently of the eyes and main body-mass.

Suppose the animal needs to use the grippers to perform various tasks, including

- pushing or pulling or twisting obstacles out of the way
- putting edible items into its mouth
- peeling or breaking objects to get at edible interiors.
- fighting off or damaging predators or competitors
- feeding or carrying young
- building shelters

What kinds of information and information-processing need to be acquired and used by such an organism, at various stages in the processes of learning and acting?

Common assumptions among AI/Robotics researchers and cognitive scientists:

1. Organisms (and robots) must use an ontology of objects with locations, orientations and properties expressed in a global cartesian coordinate system located in the environment (possibly a movable reference frame)
2. All information about the environment must be expressed in terms of sensory-motor patterns of various sorts, since sensory and motor signals are the only information the organism has access to.
   (An example of the "symbol grounding" fallacy – an old philosophical error, refuted in Kant (1781).)

But control of grippers that continuously change their shape, location, orientation, and multi-strand relationships to things grasped and obstacles, etc. (some of which also change shape at the same time) does not fit naturally in any well understood AI models.

(Except for hard-coded special cases, e.g. in factory assembly-lines.)

Towards understanding spatial structures and processes

A hand changes shape and pose as it approaches a mug in preparation for one of: lifting, pushing, pulling, rotating about a vertical axis (e.g. to get the handle into a new position), or tilting, e.g. for pouring contents (Sloman, 2005).

How is information about the structures, possibilities for change (before they occur), consequences of change, types of motion, etc. represented?

Faithful, fully detailed simulation models cannot be the answer, since their contents would also have to be understood, reasoned about and controlled: explaining ability to manipulate X by manipulation of an internal isomorphic model of X is circular:
We need a deep theory of ontologies, representation and mechanisms, for understanding

- spatial processes (including multi-strand processes, with several relationships changing concurrently),
- possibilities of processes that have not yet occurred (including branching trees of possible processes),
- possible, likely and necessary consequences of various possible processes (Piaget & *others*, 1981, 1983),
- different kinds of stuff (matter) and how, depending on shape, size and context, they respond to pushing, prodding, squeezing, stroking, twisting, pouring, folding, etc. (Sloman, 2009b).

In current AI and Robotics:

- Most commonly used forms of representation are either Fregean (e.g. using predicates, functions, relationships, quantifiers) or else numerical (e.g. using vectors of values, operations on them, etc.)
- Those are not good for reasoning at a useful level of abstraction about interactions (including discrete and continuous interactions) between changing spatial structures, shapes and relationships
- Yet pre-verbal human toddlers and many animals seem to have effective solutions for such problems in a range of cases (Sloman, 2008a).

Can we find out how they do it, and what had to evolve to enable this?

Towards a meta-ontology for Cognitive Robotics

The ontologies currently used for observation and theory-construction by most neuroscientists, psychologists, biologists, and AI/Robotics researchers cannot accommodate all of the features discussed here, seriously restricting the explanatory power of theories.

I am trying to address some ways of overcoming those limitations, in part by illustrating environmental features that have received little attention, and in part by decomposing some of the functions of a genome.

I don't have detailed theories to fill the gaps but suggest that a way to make progress would be to try to identify and organise transitions in requirements and transitions in information processing systems to meet the new requirements, in a wide range of species, including as many precursors to to current humans as possible and also different stages in the development of individual humans and other animals.

This will require interdisciplinary cooperation across a range of subjects including computer science, AI, Robotics, biology, biochemistry, ethology, neuroscience, clinical neuroscience, clinical psychology, developmental psychology, education, linguistics, social and cultural anthropology, paleontology, archeology, and philosophy.

The study would need to include transitions leading to and extending various types of meta-cognition, in parallel with (incremental) development of a cross-discipline shared ontology for describing the requirements and designs.

Perhaps this could be called The Meta-Genome, or Info-Genome Project

See also (Sloman, 2007a, 2011a)

What does biological evolution produce?

One answer is obvious: evolution produced millions of different types of organism, and billions of different individual organisms, which, in various ways, at various times, have transformed this planet.

That answer focuses only on whole organisms and on ways of classifying them.

Less obviously, in order to produce whole organisms that develop, consume resources, interact with various other things, and reproduce, evolution also had to produce

- many more different processes than there are organisms, since organisms (especially complex organisms) produce different behaviours in different circumstances
- mechanisms that make those processes possible, including
  - physical mechanisms that produce processes involving rearrangements of matter and energy
  - control mechanisms that are needed to ensure that at any time appropriate selections are made between available options for behaviour.
- designs for the mechanisms that explain how they are able to perform their functions
- implementations for the designs (often different implementations for same design)
- physical encodings for the implementations that allow multiple instances of (roughly or exactly) the same design to be produced

If we study processes that the structures can produce (some of which reproduce the structures) we find that even simple organisms need information-based control systems:

- that use information to select which competences to deploy
- that use information to control the details of the deployment of the competences (e.g. exactly how a gripper moves in order to grasp something).

Diversity of biological information processing

Viewing the processing of information as a major product of evolution, in addition to physical structures and physical behaviours, can draw attention to the diversity of contributions to biological success.

- For example, members of a species may be able at idle times to speculate about hidden explanations for observed phenomena.
- It may be that most of those speculations are of no value, but do not consume much energy or interfere with practical activities, so they do not have a serious cost.
- Yet occasionally a thought process may produce a result that is of great practical value not only to the thinker but also to conspecifics, e.g. if it leads to a new way of building shelters, or a new kind of hunting weapon, or a new technique for obtaining food based on collecting and planting seeds instead of always searching for self-seeded plants.
- But we must not assume that all thought processes are implemented in the same way (e.g. that they all involve searching for regularities in sensorimotor signals), for that may divert attention from functions requiring different mechanisms that evolution might have produced. E.g. Craik (1943) speculated that some animals can construct working "mental models" of portions of the environment and use them to devise novel solutions to practical problems.
- If we examine requirements for mental models of different aspects of the environment we can assemble clues as to what the mechanisms might be.
- So far we can neither replicate those capabilities in machines nor explain how they can be implemented in known brain mechanisms.

What we think of as simple may be deceptive

We tend to think of the information processing in certain classes of organisms as simple, e.g. insects.

But we forget that

- The information processing involved in building each insect is not yet something engineers know how to replicate.
- Many of the kinds of tasks performed by insects, e.g. ants which carry carcases larger than themselves back to the nest over terrain that is very rough, with many obstacles on the way, are far beyoned robots we know how to build.

Barbara Webb (Edinburgh University) has some interesting observations on kinds of complexity in information processing in insects.

She is interviewed here, starting about 4 minutes from the beginning, lasting about 12 minutes:

`http://www.robotspodcast.com/podcast/mp3/robots-20100618-episode54.mp3`

More specific transition-types

- Forms of representation
- Type of ontology (e.g. from somatic to exosomatic to scientific)
- Types of motive generator
- From precocial to altricial.
- Use of meta-semantic competences
- Self-monitoring of behaviours
- Self-monitoring of information processing (requires meta-semantic competences)
- Transition from behaving cooperatively to communicating
- Transition to cooperating in goal selection, and deliberation.
- Transitions in control systems that increase various kinds of freedom (Sloman, 1992).

Machines manipulating matter energy and information exist at different levels of abstraction

Many produced by biological evolution

Items on the left tend to be virtual machines (explained later) produced by biological evolution (or evolution and development).

Items on the right tend to be produced by human engineers or other machines.

All are implemented in physical and chemical mechanisms, but their important characteristics and behaviours cannot be specified using the language of the physical sciences.

Physics as we know it now is does not refer to the same entities and relationships as physics known to scientists a few centuries ago.

How many levels of (relatively) virtual machinery does physics itself require?

We cannot predict what will be added to physics in future.

For more on virtual machinery, see
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk76

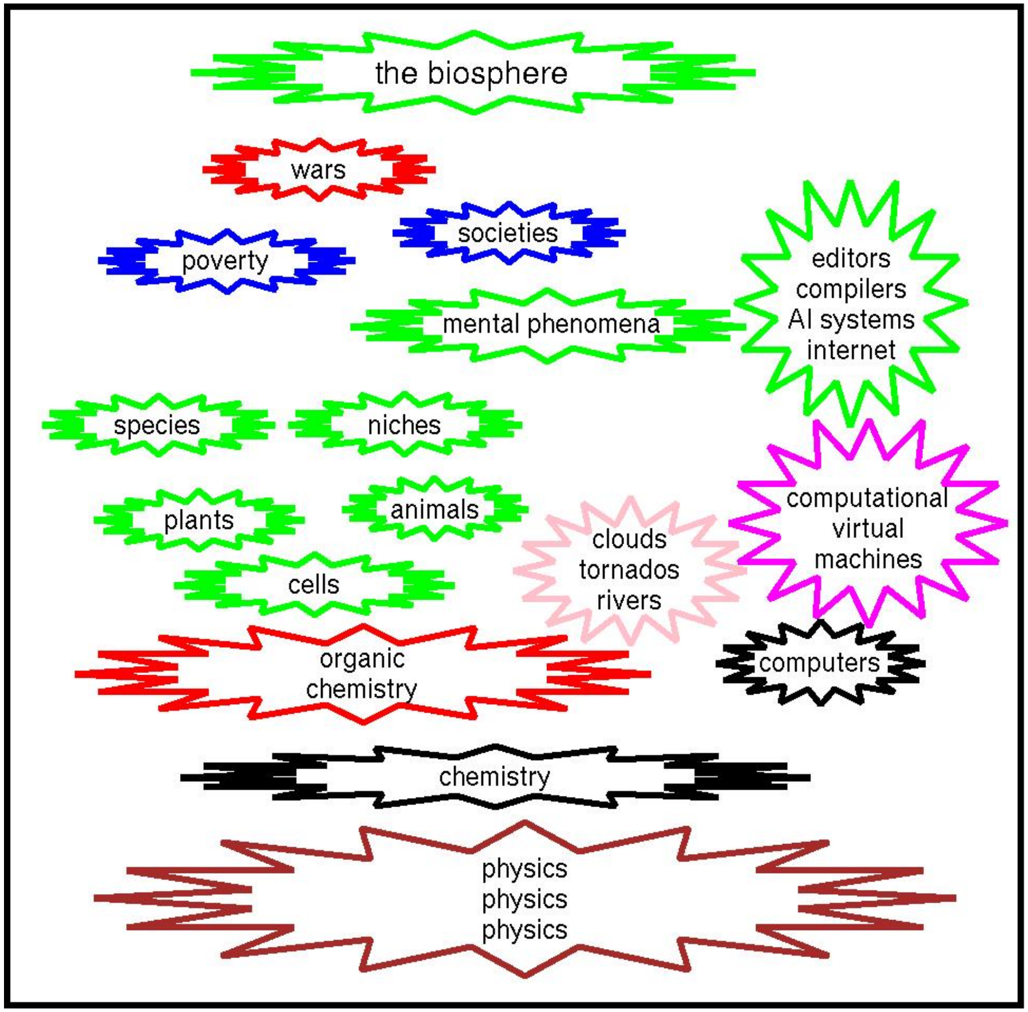

Information-processing discontinuities

This is not intended to be a complete list

- Transitions between functions (requirements, niches)
- Transitions between designs
- Transitions between implementation mechanisms
- Whether the transitions occur in evolution, or in development, or both
- Transitions in the environment and their implications
- New relationships between effects on environment of earlier members of same species and new members (e.g. effects of toys, furniture, clothes, on infants):
  - Physical effects
  - Cultural effects

NB: not all these transitions can be studied in 'toy' simulated environments for robots or simulated robots, or in tightly constrained laboratory experiments where subjects cannot use their normal competences (e.g. reaction time tests).

Some researchers in AI and some experimenters produce models or experiments in which the environment is simulated in a way that makes their task tractable. E.g.

1. 'tile-world' or '2-d Worlds' in AI experiments, or use of robots with very few sensors
2. replacement of natural tasks in which subjects are highly competent with artificial reaction-time experiments, or with experiments on the edge of normal competences.

They assume that this simplification makes no difference to the phenomenon under investigation, ignoring requirements for "ecological validity" in models and experiments.

Information-processing architectures evolved in layers

Microbes have sensors triggering processes that fairly directly produce behaviours:
– reactive architectures

Some animals can consider alternative possibilities, including branching possibilities, before acting
– deliberative architectural layers (Sloman, 2006)

Some can monitor and modulate their own information processes, using meta-semantic competences
– a meta-management architectural layer

These are very crude sub-divisions, needing much further refinement with many intermediate cases. (Some discussed in Minsky (2006).)

Different sorts of competence evolved at different stages, but did not always replace what was there before.

Often the old architecture was subsumed by the new.

Conjecture: abstract commonalities across different spatial environments caused convergent evolution of similar cognitive mechanisms (virtual machinery) in animals with different morphologies and sensorimotor systems. (E.g. apes, birds, elephants, ...)

Tradeoffs-between evolution and development

Biologists distinguish

Precocial species: individuals born physiologically advanced and competent
e.g. chicks follow hen, and peck for food; new born deer run with the herd

Altricial species: individuals born under-developed and incompetent – very dependent on parents.

Jackie Chappell convinced me in 2005 that:

It's not species but competences that should be distinguished.

Even humans have some "precocial" competences (e.g. sucking is highly developed very early).

Evolution produced very dramatic changes – over millions of years.

But it also produced some species whose individual members can make dramatic changes in their own information-processing, in only a few years –

huge changes, not in their physical shape, etc., but in their information-processing powers, and their knowledge about the world they live in.

Learning to talk, learning to read music and play musical instruments, learning to build skyscrapers, learning to make advances in theoretical physics. How?

The world is now very different from the world in which that versatility evolved: but that's why deep and rapid learning capabilities were needed.

Educators need to understand those biological mechanisms – yet no known theory is adequate: e.g. all currently implemented models of learning are simplistic by comparison.

But we can say something about the variety of routes from genome to behaviours.

Individual developmental trajectories

Routes from genome to behaviour : the direct model.

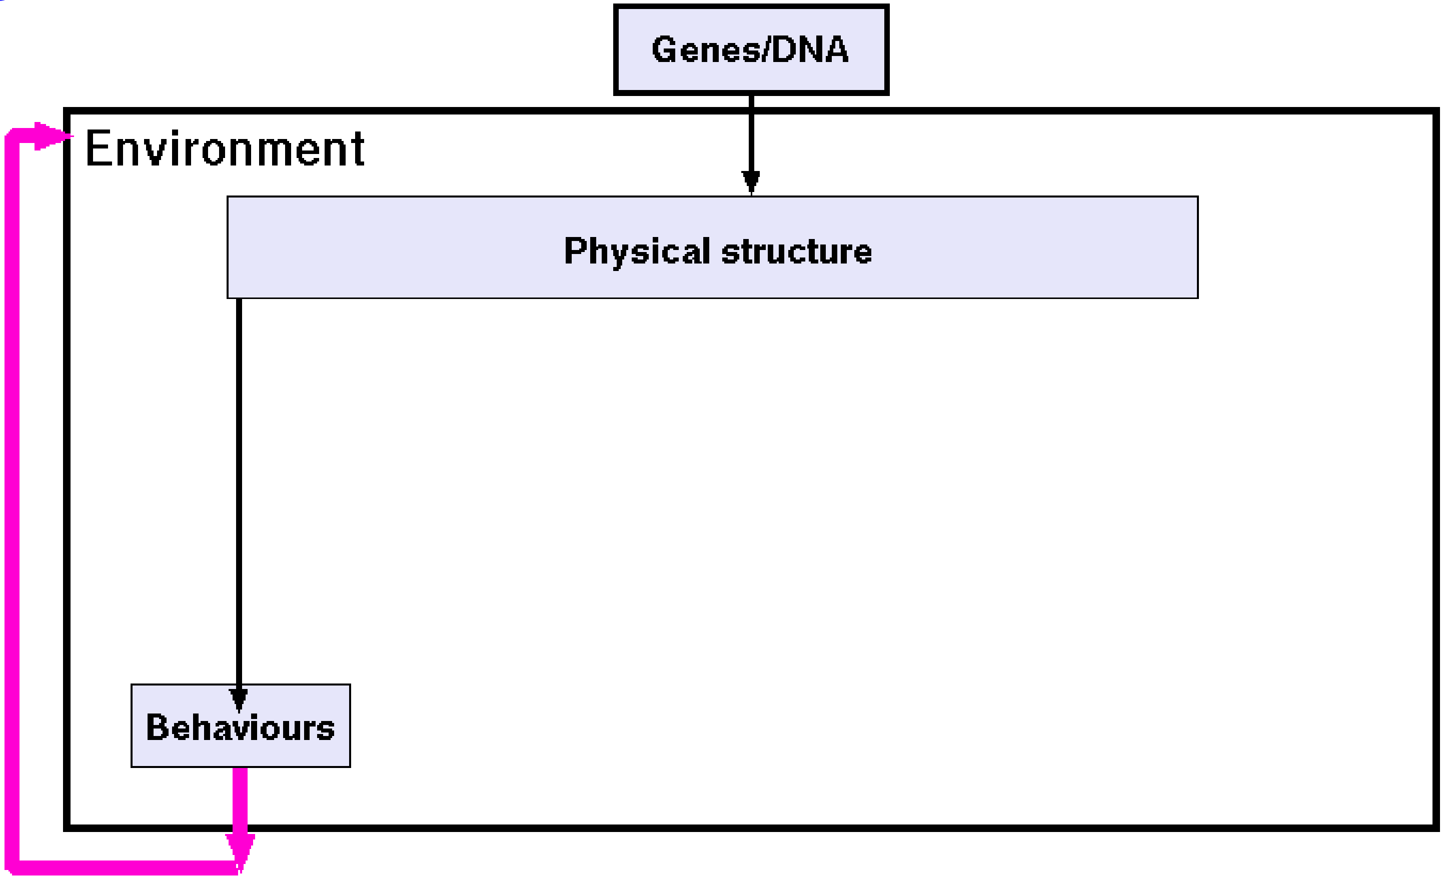

The vast majority of organisms (including micro-organisms) are like this.

Many don't live long enough to learn much – they have to make do with innate reflexes. Other organisms have more "inside the box".

(This diagram and the next few were developed with Jackie Chappell and Chris Miall)

Individual developmental trajectories

Routes from genome to behaviour : the two-stage model.

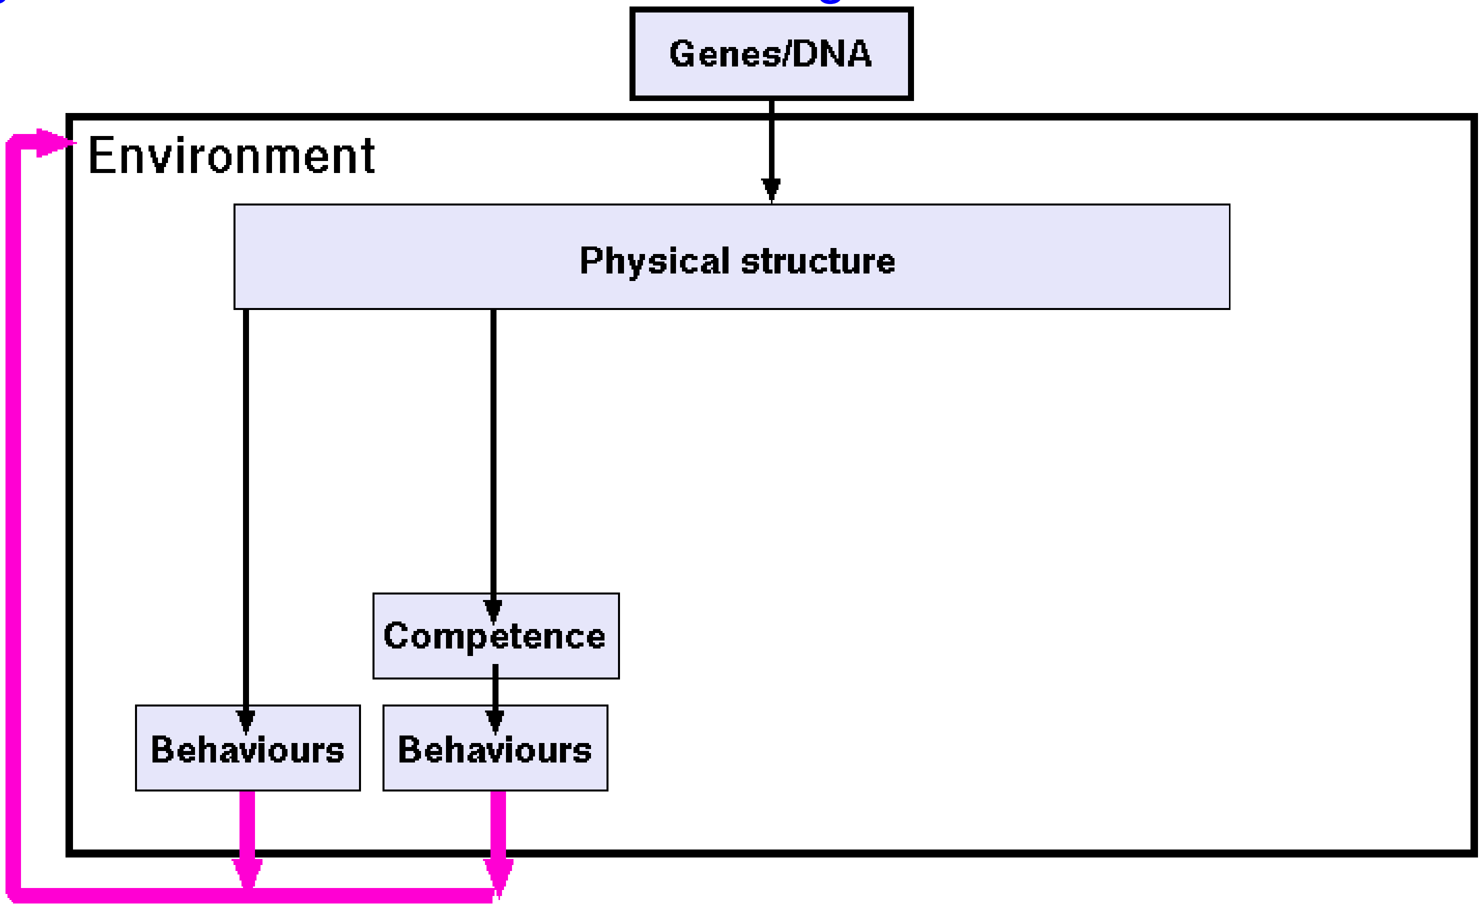

Some more complex organisms, instead of having only rigid (reflex) behaviours, also have competences that allow them to respond in fairly flexible ways to the environment: adapting behaviours to contexts.

Individual developmental trajectories

Routes from genome to behaviour : stages added by learning.

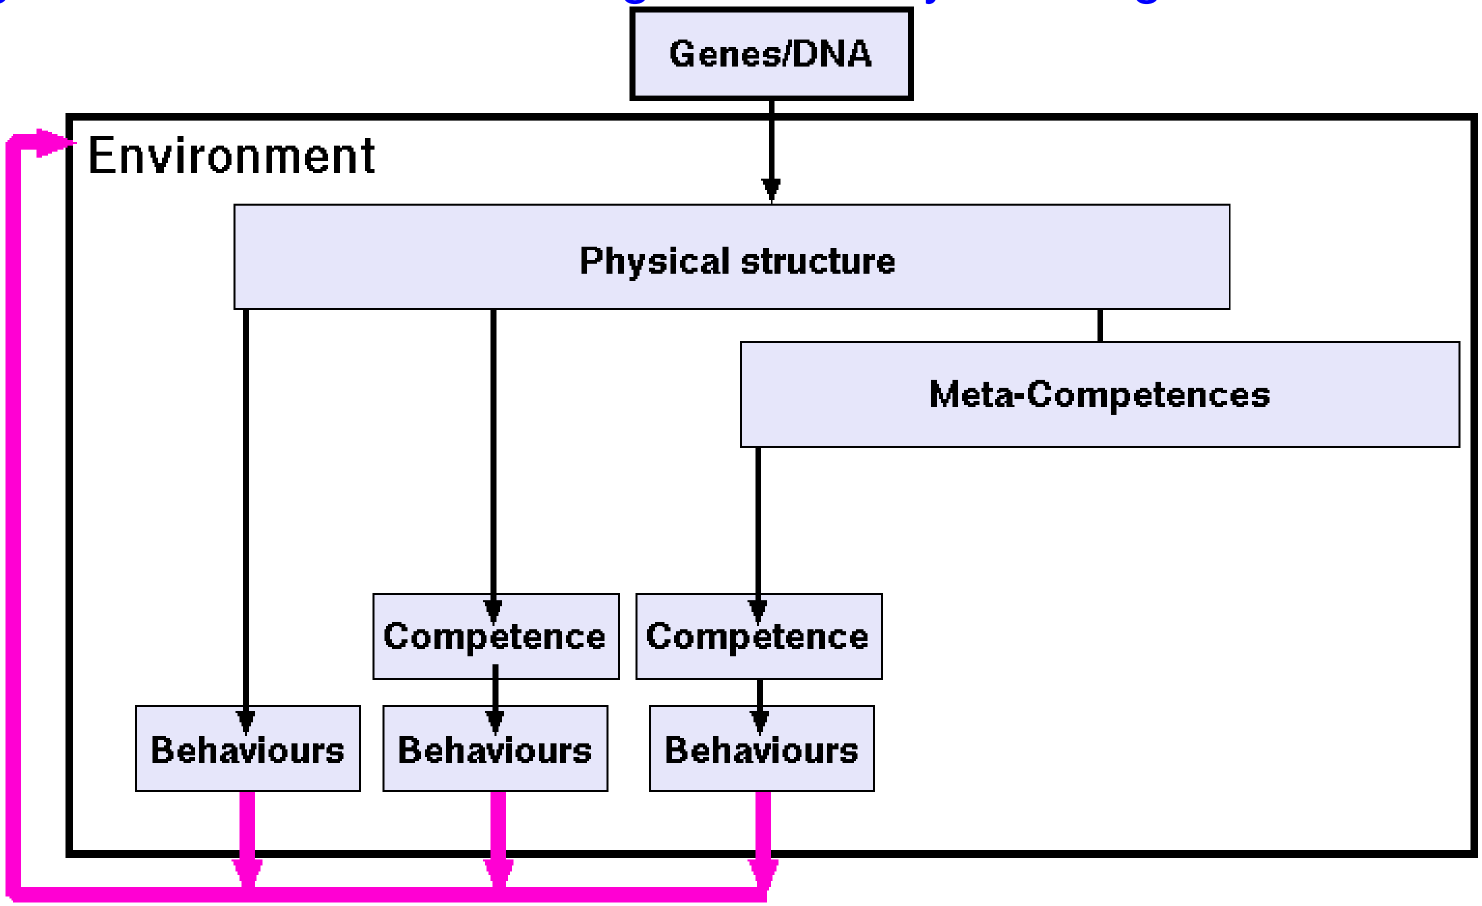

Genetically determined meta-competences allow individuals to respond to the environment by producing new types of competence, increasing flexibility and generality.

Individual developmental trajectories

Routes from genome to behaviour : the multi-stage model.

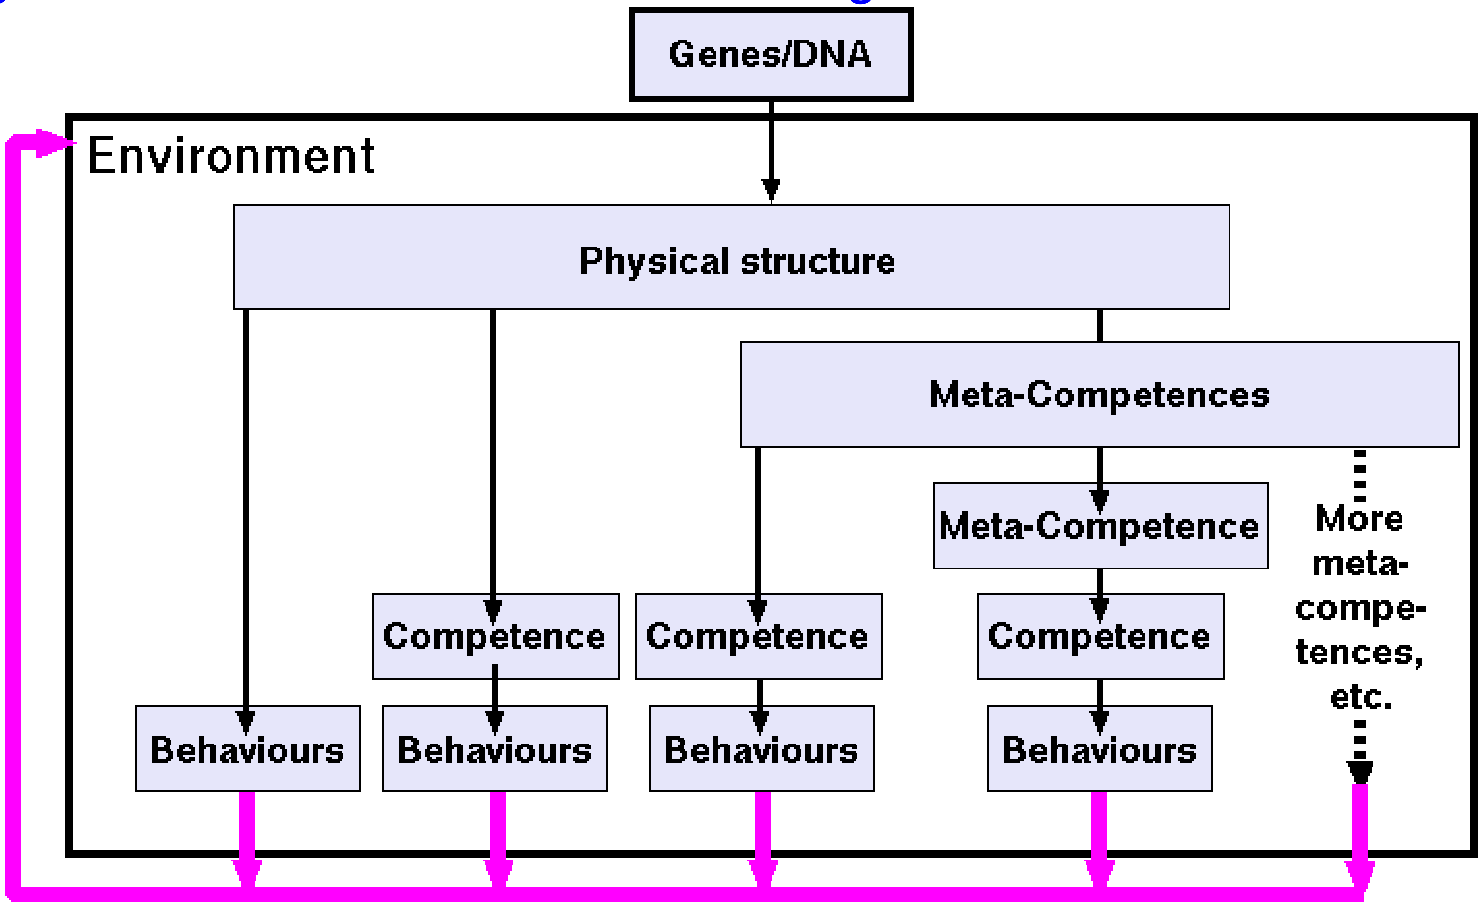

Some can also develop new meta-competences, on the basis of meta-meta competences.

Humans seem to be able to go on developing meta-meta-competences until late in life.

Mathematical competences and biological evolution

Most people think that learning mathematics requires being taught by someone who has already learnt mathematics

A QUESTION:

Who taught the first mathematicians?

ANSWER

Biological evolution and the environment we evolved in, working together.

That's the process educators are contributing to.

The bootstrapping is still continuing – and we have opportunities to push it forward.

But the opportunities are not being taken.

We need a better understanding of the problems evolution solved – in order to understand the solutions

It's hard: many problems are invisible.

Some incomplete ideas about these processes can be found in:

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler`
Why (and how) did biological evolution produce mathematicians?

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#babystuff`
Ontologies for baby animals and robots

Last words: John Maynard Smith (1920–2004)

I wish I had used the opportunity to talk to JMS more than I did when I was at Sussex University, between 1964 and 1991. He wrote, in "The Concept of Information in Biology." Maynard Smith (2000)

> "The image of development that is emerging is one of a complex hierarchy of regulatory genes, and of a signalling system that is essentially symbolic. Such a system depends on genetic information, but the way in which that information is responsible for biological form is so different from the way in which a computer program works that the analogy between them has not, I think, been particularly helpful ... "

Alas, "the way in which a computer program works" displays a lack of understanding of the diversity of forms of computation, including use of rule-based systems, event-driven programs, logic programs, the misleadingly-named object-oriented programming (use of inheritance, generic functions, multi-methods), the use of parametric polymorphism, the use of higher order functional languages (which allow programs to be inputs and outputs of other programs), and many more – often blurring the program/data distinction.

There's far more to computing systems of the last few decades than the algorithmic programming I suspect JMS had in mind, though we still have much to learn. However, I agree with the rest of the quoted sentence:

> ".... although it is a lot nearer the truth than the idea that complex dynamic systems will generate biological forms "for free".

He also wrote:

> "Given the central role that ideas drawn from a study of human communication have played, and continue to play, in biology, it is strange that so little attention has been paid to them by philosophers of biology. I think it is a topic that would reward serious study."

I've tried show some ways of doing this, while going beyond his focus on **communication**, which is only a subset of what organisms use information for. This continues the research programme promised in "The Computer Revolution In Philosophy" (1978):
`http://www.cs.bham.ac.uk/research/projects/cogaff/crp/`

(I still relish John's objection at Senate to "Propagation of error" when a degree in theology was proposed.)

References

Chappell, J., & Sloman, A. (2007). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, *3*(3), 211–239. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609)

Craik, K. (1943). *The nature of explanation*. London, New York: Cambridge University Press.

Dennett, D. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.

Gibson, E. J., & Pick, A. D. (2000). *An Ecological Approach to Perceptual Learning and Development*. New York: Oxford University Press.

Hall, W. (2006, May). Emergence and Growth of Knowledge and Diversity in Hierarchically Complex Living Systems. In *Workshop on Selection, Self-Organization and Diversity.* Katoomba, NSW, Australia: CSIRO Centre for Complex systems Science and ARC Complex Open Systems Network.

Hawkins, S. (2011). *The Mechanisms of Evolution: A Critique of the Neo-Darwinian Modern Synthesis*. Reading: Light Press.

Jablonka, E., & Lamb, M. J. (2005). *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge MA: MIT Press.

Kant, I. (1781). *Critique of pure reason*. London: Macmillan. (Translated (1929) by Norman Kemp Smith)

Lakatos, I. (1980). Falsification and the methodology of scientific research programmes. In J. Worrall & G. Currie (Eds.), *Philosophical papers, Vol I* (pp. 8–101). Cambridge: Cambridge University Press.

Maynard Smith, J. (2000, June). The Concept of Information in Biology. *Philosophy of Science*, *67*, 177–194. (2)

Maynard Smith, J., & Szathmáry, E. (1995). *The Major Transitions in Evolution*. Oxford, England:: Oxford University Press.

McCarthy, J. (2008). The well-designed child. *Artificial Intelligence*, *172*(18), 2003-2014.

Minsky, M. L. (1963). Steps towards artificial intelligence. In E. Feigenbaum & J. Feldman (Eds.), *Computers and thought* (pp. 406–450). New York: McGraw-Hill.

Minsky, M. L. (1992). Future of AI Technology. *Toshiba Review*, *47*(7).

Minsky, M. L. (2006). *The Emotion Machine*. New York: Pantheon.

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

Peterson, D. (Ed.). (1996). *Forms of representation: an interdisciplinary theme for cognitive science*. Exeter, U.K.: Intellect Books.

Piaget, J., & *others*. (1981). *Possibility and Necessity Vol 1. The role of possibility in cognitive development*. Minneapolis: U. of Minnesota Press. (Tr. by Helga Feider from French in 1987)

Piaget, J., & *others*. (1983). *Possibility and Necessity Vol 2. The role of necessity in cognitive development*. Minneapolis: U. of Minnesota Press. (Tr. by Helga Feider from French in 1987)

Popper, K. (1972). *Objective Knowledge*. Oxford: Oxford University Press.

Rochat, P. (2001). *The Infant's World*. Cambridge, MA: Harvard University Press.

Shanahan, M. (2010). *Embodiment and the inner life: Cognition and Consciousness in the Space of Possible Minds*. Oxford: OUP.

Sloman, A. (1971). Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In *Proc 2nd ijcai* (pp. 209–226). London: William Kaufmann. (http://www.cs.bham.ac.uk/research/cogaff/04.html#200407)

Sloman, A. (1978). *The computer revolution in philosophy*. Hassocks, Sussex: Harvester Press (and Humanities Press).

Sloman, A. (1979). The primacy of non-communicative language. In M. MacCafferty & K. Gray (Eds.), *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979* (pp. 1–15). London: Aslib. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43)

Sloman, A. (1992). How to Dispose of the Free-Will Issue. *AISB Quarterly*, *82*, 31–32. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#8)

Sloman, A. (1993). The mind as a control system. In C. Hookway & D. Peterson (Eds.), *Philosophy and the cognitive sciences* (pp. 69–110). Cambridge, UK: Cambridge University Press. (http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#18)

Sloman, A. (2000). Interacting trajectories in design space and niche space: A philosopher speculates about evolution. In *et al.*. M.Schoenauer (Ed.), *Parallel problem solving from nature – ppsn vi* (pp. 3–16). Berlin: Springer-Verlag.

Sloman, A. (2005). *Perception of structure: Anyone Interested?* (Research Note No. COSY-PR-0507). Birmingham, UK.

Sloman, A. (2006, May). *Requirements for a Fully Deliberative Architecture (Or component of an architecture)* (Research Note No. COSY-DP-0604). Birmingham, UK: School of Computer Science, University of Birmingham.

Sloman, A. (2007a). *A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots* (No. COSY-DP-0701). (Contribution to euCognition wiki)

Sloman, A. (2007b). *Why symbol-grounding is both impossible and unnecessary, and why theory-tethering is more powerful anyway.* (Research Note No. COSY-PR-0705). Birmingham, UK. (http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models)

Sloman, A. (2008a). *Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)?* (Research Note No. COSY-PR-0702). Birmingham, UK. (http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0702)

Sloman, A. (2008b, July). Kantian Philosophy of Mathematics and Young Robots. In S. Autexier, J. Campbell, J. Rubio, V. Sorge, M. Suzuki, & F. Wiedijk (Eds.), *Intelligent Computer Mathematics* (p. 558-573). Berlin/Heidelberg: Springer.

Sloman, A. (2009a). Architecture-Based Motivation vs Reward-Based Motivation. *Newsletter on Philosophy and Computers*, *09*(1), 10–13.

Sloman, A. (2009b). *Ontologies for baby animals and robots. From "baby stuff" to the world of adult science: Developmental AI from a Kantian viewpoint.* University of Birmingham, School of Computer Science. (Online tutorial presentation)

Sloman, A. (2009c). Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. In B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, & K. Doya (Eds.), *Creating Brain-like Intelligence* (pp. 248–277). Berlin: Springer-Verlag.

Sloman, A. (2011a). Varieties of Meta-cognition in Natural and Artificial Systems. In M. T. Cox & A. Raja (Eds.), *Metareasoning: Thinking about thinking.* Cambridge, MA: MIT Press.

Sloman, A. (2011b). What's information, for an organism or intelligent machine? How can a machine or organism mean? In G. Dodig-Crnkovic & M. Burgin (Eds.), *Information and Computation* (pp. 393–438). New Jersey: World Scientific.

Sloman, A., & Chappell, J. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings IJCAI'05* (pp. 1187–1192). Edinburgh: IJCAI. (http://www.cs.bham.ac.uk/research/cogaff/05.html#200502)

Sloman, A., Chappell, J., & CoSyTeam, T. (2006, June). How an animal or robot with 3-D manipulation skills experiences the world. In *The tenth annual meeting of the Association for the Scientific Study of Consciousness, Oxford.* Internet: ASSC. (Poster for ASSC10, Oxford June 2006. Also at ASSC10 Eprints Archive: http://eprints.assc.caltech.edu/112/)

Turing, A. (1950). Computing machinery and intelligence. *Mind*, *59*, 433–460. ((reprinted in E.A. Feigenbaum and J. Feldman (eds) *Computers and Thought* McGraw-Hill, New York, 1963, 11–35))