

**IJCAI 2016 TUTORIAL**  
**International Joint Conference on AI**  
**New York, July 2016**

---

**Note: 25 Aug 2021: More recent developments  
(likely to be revised)**

Since late 2020, consideration of largely unnoticed facts about developments that occur in eggs, leading to production of highly competent hatchlings, such as the avocets in this 35sec video clip from BBC Springwatch Episode 9, 2021,

<https://www.cs.bham.ac.uk/research/projects/cogaff/movies/avocets/avocet-hatchlings.mp4>

triggered a complex collection of new ideas about the roles of chemistry in production of deep forms of spatial intelligence in many organisms. I have been giving talks about this during 2021, and the next one is scheduled for 15th Sept 2021, for which notes are in preparation here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-morcom.html>

Mike Levin at Tufts University has agreed to respond.

There are still deep gaps in the theory. I hope other researchers with relevant knowledge of biology, chemistry, and processes of development of spatially intelligent animals, can contribute relevant information and hypotheses, significantly extending the ideas presented at the IJCAI tutorial in 2016,

---

**Note: 16 Feb 2017**

**This web page started life as an advertisement for the tutorial, presented at IJCAI in 2016, in New York.**

## **Tutorial T24**

**If Turing had lived longer, how might he have investigated what AI and Philosophy can learn from evolved information processing systems?**

**This tutorial presents tentative, partial answers developed in this project, triggered by the Turing centenary year:**

**The Meta-Morphogenesis Project  
(Or Self-informing Universe Project)**

**Sunday July 10th, 2016**

**Including homage to John McCarthy and Marvin Minsky, two of the founders of AI, recently deceased, both interested in connections between AI and philosophy.**

**Presenter: Aaron Sloman**  
**School of Computer Science, University of Birmingham, UK**

<http://www.cs.bham.ac.uk/~axs>  
(More information about presenter [below.](#))

**Tutorial web site:**

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-tut-ijcai-2016.html>  
Abbreviated version: <http://goo.gl/8Bfizj>

---

## **BACKGROUND TO THIS TUTORIAL**

**Homage to John McCarthy (1927-2011) and Marvin Minsky (1927-2016)**

[https://en.wikipedia.org/wiki/John\\_McCarthy\\_\(computer\\_scientist\)](https://en.wikipedia.org/wiki/John_McCarthy_(computer_scientist))

[https://en.wikipedia.org/wiki/Marvin\\_Minsky](https://en.wikipedia.org/wiki/Marvin_Minsky)



Images courtesy of Wikimedia.

There are many online tributes to both. My tribute concerns their contributions to philosophy.

From 1969, when [Max Clowes](#) introduced me to AI, the writings of McCarthy and Minsky convinced me of the deep relevance of AI to philosophy, even when I disagreed with them.

Both still have many papers freely available on their web sites:

<http://www-formal.stanford.edu/jmc/>

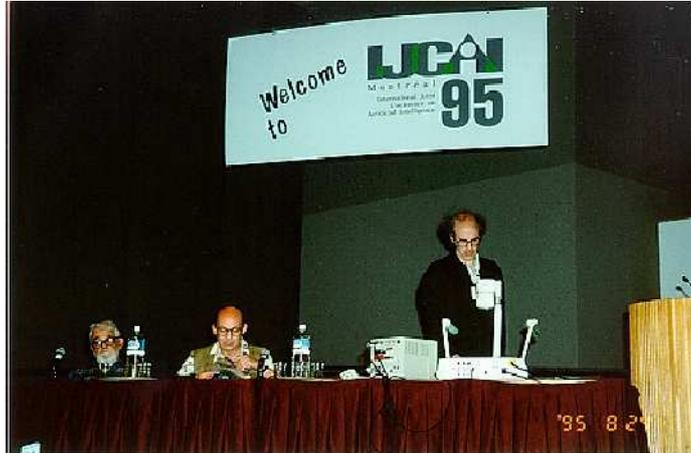
<http://web.media.mit.edu/~minsky>

(McCarthy's presentations were not properly formatted for landscape:  
fetch the .tex files, add [landscape] to the headers and run pdflatex,)

Both McCarthy and Minsky were primarily interested in AI as [science](#) rather than [engineering](#). They focused on phenomena that are hard to explain and attempted to characterise requirements, forms of representation, forms of reasoning, architectures, mechanisms etc., relevant to finding explanations. This leads to work on hard, long term theoretical and practical problems rather than work on making machines do something useful in the near future. Minsky, in particular, lamented the pressures from academic institutes and research funders pushing very bright researchers into work on short term practical problems, or problems for which new technologies make demonstrable results reachable fairly quickly, rather than working on hard, deep, scientific problems, on which

progress is likely to be slow and unpredictable.

-----  
**Below:** Both at the "Philosophical Encounter" Session with AS at IJCAI 1995, Montreal



Many thanks to Takashi Gomi, at Applied AI Systems Inc, who took the picture.  
(I have lost the original, higher resolution, image file.)

**An audio recording of the 2hour session was made by  
Audio Archives International Inc. Does anyone have the recording?**

Personal tributes:

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-jmc-aisb.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-minsky-tribute.html>

---

## NOTE

There are video recordings of a number of interviews, lectures and tutorials on various aspects of the Meta-Morphogenesis project available here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/movies>

---

## CONTENTS

- [Homage to John McCarthy and Marvin Minsky \(above\)](#)
- [OVERVIEW + Summary \(Below\)](#)
- [NOTE on the "AI Singularity"](#)
- [The Meta-Morphogenesis \(Self-Informing Universe\) Project](#)
- [How it started](#)
- [The need to understand requirements](#)
- [No behavioural test can specify requirements for intelligence](#)
- [Figure: Evolutionary Steps in Information processing](#)
- [Evolution of sensing and perceiving](#)
- [Types of biological perception mechanism](#)
- [From matter-manipulation to information-manipulation](#)

- [Figure: Room view](#)
- [Meta-reasoning](#)
- [Mathematical competences](#)
- [EVO-DEVO Issues](#)
- [Figure Evo-Devo](#)
- [Implications for evaluation of theories of intelligence](#)
- [Evolution as a blind mathematician](#)
- [Inspiration from Alan Turing](#)
- [Evolution of mathematical abilities](#)
- [Figure: Blocks](#)
- [The need for construction kits](#)
- [What is information?](#)
- [The fundamental use of information is for control](#)
- [Evolutionary transitions from microbes to intelligent animals](#)
- [What don't we understand yet](#)
- [Criticisms of some current fashions](#)
- [Rich internal languages must have evolved before languages for communication](#)
- [The problem of specifying the functions of vision in intelligent systems](#)
- [Varieties of motivation](#)
- [Architectures](#)
- [Virtual machines and consciousness](#)
- [Sample References](#)
- [Original draft proposal for half-day tutorial](#)

[Back to contents](#)

---

## OVERVIEW

Bullet-points used in my presentation at IJCAI on 10th July 2016 have been expanded and reorganised here. I have tried to make this version of the tutorial notes fairly self-contained, but with many external links, for those who would like greater breadth and depth. This revised version starts with a high-level summary of aspects of the M-M project.

### SUMMARY

There are many deep, but not widely recognized, gaps between natural intelligence (NI) (in humans and non-humans) and products of AI research and development. Examples of such gaps are collected in a separate document:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-ai-gaps.html>

The gaps may not matter for narrowly focused AI applications -- AI as engineering. However they do matter for AI construed as the science of intelligence, aiming to use computational theories to answer deep questions about the variety of forms of intelligence, including intelligence in humans and other animals.

The emphasis on uses of information has begun to draw attention to the diversity of forms of representation, i.e. languages produced by evolution (or by individual development guided partly by the environment and partly by the genome), including languages for purely **internal** uses in perception, intending, wondering, noticing, discovering, planning, deciding, carrying out plans, and many more. If these occur in many intelligent non-human animals, that must **completely** revise our view of the nature of language (illustrated below and in [\[Sloman\(Vision\)\]](#)).

**Epigenesis:** Attending to previously unnoticed transitions in cognitive development of individuals in intelligent human and non-human species may also provide clues (e.g. use of topological information by pre-verbal human toddlers and other animals)[\[Chappell Sloman 2007\]](#), [\[Karmiloff-Smith 1992\]](#).

Like McCarthy and Minsky, I focus more on AI as **science** and **philosophy** than AI as **engineering**, though the interests overlap: many aspects of AI as engineering depend on good science and philosophy. In particular, researchers in AI (and cognitive science) who know nothing about the work of Kant and other great philosophers risk missing some of the deepest features of minds, language, and thought that need to be explained and modelled.

The reverse is also true: philosophers with shallow understanding of the science and engineering issues in computing and AI, including what we have learnt about varieties of virtual machinery since Turing died, will produce shallow philosophical theories of mind, language, science, mathematics, etc.

Many researchers remain mystified, or even mystical, about mental phenomena because their education has not introduced them to the required types of explanatory mechanism -- mechanisms capable of filling the so-called "Explanatory Gap" (pointed out by Darwin's admirer T.H.Huxley [\[Huxley 1866/1872\]](#) and repeatedly re-discovered, and re-labelled, since then)[\[SEP "Consciousness"\]](#). (Huxley toned down his wording in the 1872 edition.)

### **NOTE on the AI Singularity"**

I shall ignore the so-called "AI Singularity", about which much has been written that is ill-informed about the science, over-optimistic, or over-pessimistic, and to me mostly uninteresting. (I think good science fiction writers generally produce deeper and more interesting imagined futures.) I have a brief note on why I (mostly) ignore the topic [here](#). The main reason is that I want to focus on a collection of deep and difficult problems that almost all researchers (including my heroes) have ignored, though Immanuel Kant's writings drew my attention to some of them, e.g. issues about the nature of mathematics and causation, and how intelligent individuals think about and use them. I think Jean Piaget also partly understood some of the problems, while lacking explanatory theories (as shown by his last two books on Possibility and Necessity, published posthumously [\[Piaget 1981-1983\]](#)), though much research in developmental psychology, cognitive science and neuroscience ignores the problems, under pressure from academic evaluation procedures that emphasise statistics-based publication requirements.

This tutorial presents a subset of such unsolved or only partially solved problems, some of which I have been thinking about since before I knew anything about AI, such as the problem of explaining what kinds of mechanisms could have led to the mathematical discoveries made by ancient mathematicians, including Euclid and Archimedes, [\[Sloman 1962\]](#). My ideas build

on those in [\[Kant 1781\]](#). It is also arguable that biological evolution made use of many mathematical discoveries -- some discussed [below](#).

Other problems include characterising the functions and mechanisms of human and animal vision, including its role in mathematical discovery. There are many unexplained features of natural vision, including the speed at which humans can switch between visual scenes, e.g. in some television adverts that switch scenes approximately every second. Some of these abilities can be probed in laboratory experiments. Others, like the role of vision in mathematical discoveries, are much harder to study empirically.

Current AI vision systems need considerable training on image structures associated with different visual scenes. In contrast a human who is not a botanist can walk around a botanical centre enjoying seeing rich and varied, mostly unfamiliar collections of plants.

There are also many aspects of the development of young humans and other intelligent species about which there are deep problems. E.g., as explained below, although human children seem to be **learning** the language of their environment from examples of communications by older speakers (or signers), there is suggestive evidence that much of what appears to be learning from experience is actually **collaborative creation** of a language, where the youngest collaborators are usually in a minority.

Other gaps in current AI are concerned with affective states: including likes, dislikes, enjoyment, preferences, values, attitudes, moods, emotions, ideals, ambitions, aesthetic reactions, and finding things funny. Most of the research in psychology, neuroscience and AI that attempts to address such issues seems to me to barely scratch the surface, compared with what is understood and used by great novelists, playwrights, poets, comedians, and some counsellors. I shall say very little here about such matters though our Cognition and Affect project [\[Sloman et al 1991--\]](#) has addressed some of the issues, and a former student, Luc Beaudoin has a growing web-site [\[Beaudoin, CogZest\]](#) expanding the discussion.

Some of the issues were also addressed in Marvin Minsky's last two books [\[Minsky 1987\]](#) and [\[Minsky 2006\]](#), though his interest was more human-centred than mine, and as far as I know he (like McCarthy) mostly ignored ancient human mathematical competences, some of which are shared with other animals and begin to develop in pre-verbal toddlers.

## **The Meta-Morphogenesis Project (Self-Informing Universe Project)**

The Meta-Morphogenesis (M-M) project attempts to address these hard unsolved problems by exploring a tentative answer to the question: "*What would Alan Turing have done if he had not died two years after publication of his paper 'The Chemical Basis of Morphogenesis' [Turing, 1952] and had lived three or four more decades?*" My tentative answer has four parts (not claimed to be a complete list!):

1. He would have investigated the development of varieties of *information processing* between the earliest organisms, or pre-life molecules, and current biological organisms. (This investigation differs from but complements evolutionary investigations that many scientists are engaged in, e.g. changes in: physical forms, observable behaviours, forms of locomotion, types of habitat -- e.g. water, land, air, deserts, polar regions, deep sea vents, etc. -- types of nutrient, and genome details. These are all connected with but

different from changes in information processing.)

2. Some previously unrecognized intermediate forms of biological information processing may be important clues about hitherto unrecognized forms of information processing in the brains of known intelligent species, including humans. Human brains are so complex that we may have missed crucial features of their operation that would be easier to identify in less advanced evolutionary products.
3. Those mechanisms may play essential roles in competences of humans and other animals that current AI systems do not replicate and don't even seem to be close to replicating, including the competences of many intelligent non-human species (e.g. squirrels and weaver birds), the competences of pre-verbal human toddlers [[Sloman-Toddler](#)], and the abilities that enabled ancient mathematicians such as Euclid and Archimedes to make deep discoveries, especially in geometry, that have proved to be of profound importance for science and technology, but which used modes of discovery and reasoning that are beyond the reach of current artificial mathematical reasoners, and (as far as I know) cannot be explained by current models and theories in neuroscience. For example, it is provably impossible to use the constructions permitted in Euclidean geometry to trisect an arbitrary angle, yet ancient mathematicians discovered a simple extension to Euclidean geometry that makes it easy to trisect any angle, as explained in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/trisect.html>
4. Besides the actual mechanisms used by living organisms there are various types of "construction kit", some also produced by evolution, that are used in production of biological information-processing systems. It is possible that some of those construction kits can produce types of information processing systems used by brains, but still unrecognized by human scientists and engineers. Evolved construction kits are mentioned briefly [below](#).

Inspired by these reflections on what Turing might have done had he lived several more decades, the M-M project has four strands (at present):

1. Systematic survey of changes in forms of biological information processing since [the earliest life forms](#). (This is a huge and multi-faceted project!)
2. Search for evidence that human brains include previously unnoticed examples of such intermediate forms of information processing.
3. Locate competences of humans and other animals that cannot be explained by currently understood information processing mechanisms, and investigate whether one or more of the previously unnoticed intermediate forms can help to fill the gaps.
4. Identify types of construction kit used in biological evolution and epigenetic mechanisms and investigate whether they may be relevant to filling gaps in the current catalogue of available computational mechanisms.

This tutorial introduces the problems and some conjectured partial answers.

### **How it started**

The project began towards the end of 2011, triggered by an invitation to contribute to a volume celebrating the centenary of Alan Turing's birth:

*Alan Turing: His Work and Impact*, Editors: S. Barry Cooper, J. van Leeuwen. More detailed information on the book is available [here](#). The book includes the publication announcing the M-M project as a conjectured answer to the question "What would Turing have done in the next few decades if he had not died two years after his morphogenesis paper was published?". The paper starts on page 849: "Meta-Morphogenesis: Evolution of Information-Processing

Machinery", followed up with a growing collection of papers on this web site: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html> (My three earlier contributions in that volume are also relevant.)

Barry Cooper died prematurely in October 2015. A very short tribute is [here](#).

Information processing is the basis of all forms of control in living organisms, including simple single celled organisms [[Ganti, life](#)]. There is a huge variety of forms of biological information-processing, indicated below. It is an open question whether models and mechanisms so far developed in Cognitive Science and neuroscience, or in practical applications of AI, explain or replicate, or can in future explain or replicate, natural intelligence -- e.g. the intelligence of nest building birds, squirrels, elephants, human toddlers, or ancient mathematicians such as Euclid and Archimedes?

In order to address that question we need to have accurate and detailed characterisations of the competences to be replicated -- what engineers would call "requirements specifications".

The target competences of AI research projects are often specified in a manner that ignores important details of natural intelligence, including the uses of visual perception in many animals. I'll give examples below, especially examples concerned with perception of *possibilities* and discovery of kinds of *impossibility* and *necessity*, e.g. in geometry and topology.

Many researchers ignore these complications and focus on getting machines to learn regularities and probabilities: a much shallower task that excludes most of the deep human discoveries in mathematics and science. (A more complete discussion would also need to include Meta-Requirements of various sorts [[Sloman, Vernon 2007](#)]).

I'll discuss the potential for AI to learn from other disciplines, including Biology and Philosophy, and the potential of other disciplines, including Biology, Philosophy and Psychology, to learn from AI.

If models and mechanisms so far developed and used in AI are insufficient, what would be sufficient? Can suitably designed, but far more complex, computer-based machines (physical and virtual machines) match the products of evolution? Or will entirely new information processing mechanisms be required, such as machines based on chemical information processing (the basis of much of life).

For example, could some new type of computer support artificial mathematicians able to make the sorts of discoveries in geometry and topology made by ancient mathematicians, but not yet replicated in AI reasoners? (Examples of such discoveries are given below.) The requirements for such systems are still far from clear.

### **The need to understand requirements**

Identifying those ill-understood requirements is a major task for AI: currently assumed requirements leave out too much of the functionality of biological intelligence (e.g. intelligence of squirrels, crows, elephants and pre-verbal human toddlers).

Moreover, the aspects of natural intelligence that researchers aim to replicate, model or explain are often characterised in grossly over-simplified terms. For example most AI researchers studying human natural language competences completely ignore human sign languages used in

many countries, which have many features that are very different from spoken and written languages, yet suffice for rich forms of linguistic interaction. (There are over three hundred sign languages in use.)

There is clear evidence that the (normal) human genome not only provides newborn humans with potential to learn any human spoken or written language, but also any sign language. The potential may go unrealised either because of lack of social opportunities or because of some peripheral genetic defect, e.g. deafness or blindness.

Another type of example, ignored by most researchers in human and machine vision, is the use of vision or visual imagination in various kinds of mathematical discovery, e.g. discoveries in geometry and topology. Any theory or design specification that does not account for such capabilities is an inadequate theory of human vision.

The situation is worse than that insofar as human visual mathematical competences have precursors and analogues in other animals that are able to see possibilities and impossibilities -- i.e. affordances (generalising Gibson's notion of affordance, as explained below).

So instead of asking for a test or set of tests [for intelligence](#) we should try to specify requirements and tests for as many different [types](#) of intelligence as are found in nature, and also for future possible types of natural (e.g. evolved) or artificial (designed) intelligence.

That would include tests for a machine having human-like or squirrel-like intelligence, among others. Moreover, tests for human intelligence would have to vary according to stage of development, physical environment, and cultural opportunities and norms. Some examples of toddler intelligence, including proto-mathematical competences can be found in [\[Sloman-Toddler\]](#).

But no collection of tests is a substitute for a deep theory from which the tests could be derived. Any collection of tests could be passed by a "fake" system designed only to perform in those test situations and with no other competences.

From an engineering standpoint the question should be posed as: "What precisely are the requirements to be satisfied, for various kinds of intelligence?" But the answer will not be a set of behavioural tests to be passed by AI systems.

[A better answer would provide a general and deep specification of the competences that need to be provided by the mechanisms used, and how those competences can vary across developmental stages, or in different geographical or cultural environments.](#)

### **No behavioural test can specify requirements for intelligence**

Trying to propose a single behavioural test for intelligence is pointless because there are so many varieties of natural intelligence for which different tests would be required. (That's because the concept "intelligence" is polymorphous, like "efficiency" and "consciousness", as explained in [\[Sloman-Poly\]](#).)

Trying to create a behavioural test for something being intelligent is as pointless as trying to create a behavioural test for something being efficient! There certainly are tests for efficiency but different tests are needed for different kinds of efficiency: e.g. compare efficiency in a lawn-mower, a manufacturing process, a steam turbine, a proof, a form of medical treatment, a police force or

different sorts of computer program. Likewise there are tests for various kinds of intelligence required by different sorts of animals or humans engaged in different sorts of activity.

Likewise, it would be pointless to try to specify a test for something to be a compiler, by specifying a set of programs that it should be able to compile for a particular set of target machine. Such a test would not be adequate for compilers for other languages, and other target machines. Instead of a behavioural test we need a generic specification of requirements, and meta-requirements.

Rather than a test to decide whether an **individual** is intelligent, we need a collection of tests to decide whether a **theory** of intelligence has adequate explanatory power ([Sloman \[2009\]](#)). Compare a theory of compilers and their requirements, or theory of chemical compounds and reactions.

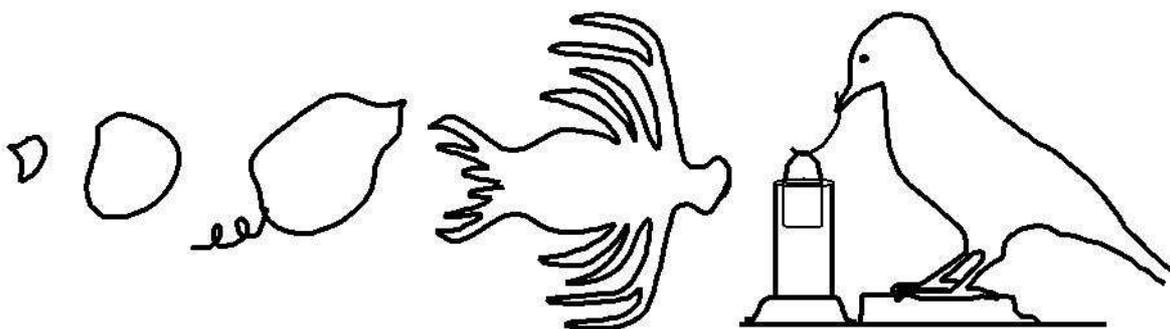
A good theory of chemistry should cover all the possible varieties of chemicals and chemical reactions. Likewise a good theory of intelligence should cover all the varieties of intelligence in animals of various sorts, and the kinds of intelligence at work (visibly and invisibly) at various stages of development of each species. (Similar points are made about varieties of consciousness in [\[Sloman-Poly\]](#), and varieties of motivation [below](#).)

Such a theory should support research and development on many particular types of intelligence, with particular sub-classes of competences.

Current ideas about what intelligent organisms can do are too restricted, e.g. with too much emphasis on learning *regularities*, which I'll contrast with learning about *possibilities* and *impossibilities* (e.g. in mathematical discovery).

We need better **meta-theories** about requirements for explanatory theories.

**Figure: Evolutionary Steps in Information processing**



"Intelligence" is a **polymorphous** concept with many different manifestations and functions in different organisms and in future machines ([\[Sloman-Poly\]](#)).

### **Evolution of sensing and perceiving**

For example, an aspect of intelligence is perception: acquiring potentially relevant information from the environment.

Forms of biological perception differ in the features of the environment providing the information, the content of the information provided, how the information is acquired, how the information is interpreted (e.g. used to derive new information, or to control actions, or to influence decisions, etc.), and possibly also differ in whether they can be replicated or emulated in computer-based machinery.

Varieties of perception also differ in the physical devices/sensors acquiring the information (modality), whether the information is used immediately and discarded ([online](#) intelligence) or stored for possible (varied) uses later on ([offline](#) intelligence).

The type of information acquired can also change from moment to moment, for example in a bird that can fly high above trees, fly through trees, and move around on a branch, or on the ground below -- all using pervasive constantly changing information.

The types of perceptual information acquired by organisms are not all the same. Perceptual information includes information about physical aspects of the environment. But that includes many sub-types of information, e.g. what surfaces are visible, what the surfaces belong to, what motions are occurring, what effects they are having, and how those features of the environment may be useful or obstructive to potential actions of the perceiver. (Gibson's positive and negative "affordances" [[Gibson 1979](#)]).

Predecessors of ancient mathematicians like Euclid must have found ways of extending information that could be acquired from the environment, by relating perception and reasoning: for example reasoning about possible changes of shape, changes of spatial relationships, changes of size, etc. and the implications of and constraints on such changes. (These extend Gibson's affordances [[Sloman\(Vision\)](#)].)

Each of these sub-types of perceptual information, and the types of use of perceptual information, can differ across species, across individual members of a species, and also across the life of an individual that develops physically and cognitively, though evolution seems to have found commonalities that allow re-use of designs, by abstraction and parametrisation, discussed below.

Many of these changes are very subtle and difficult to study. They vary in the detailed competences and the order of development, and can also vary across environments and cultures -- defeating most attempts to produce generalisations about child development. I'll offer a schema for accommodating that variation in the section on "Evo-Devo issues" [below](#).

The information may be about other agents, including their physical structures and motions, their locations, and also non-physical aspects, such as what the agents can or cannot do, their intentions, what they know or could come to know, and many more.

-----  
[Back to contents](#)  
-----

### **Types of biological perception mechanism**

Mechanisms for acquiring information about the environment have varied enormously across evolutionary time scales, from the very simplest organisms (e.g. single-celled organisms) through many intermediate forms with varying types of complexity.

These evolutionary changes depend in part on changes in the environments inhabited, and also changes in the physical forms, the physical sensor capabilities, the biological needs, the information-processing capabilities, and the mechanisms used for implementing those capabilities -- e.g. in organisms with and without brains, or with brains of different sorts.

Major divergences of functions of perception (and related forms of information processing) concern divergence between organisms that can and cannot move, whether their only form of motion is growth (e.g. in many plants) that can and cannot change their spatial configuration (e.g. with mobile limbs), that have different biological needs at different times.

For organisms that can move there is a difference between using information for online control e.g. for steering, and for offline control e.g. use in planning before acting, or analysing results of action for future use. This is discussed further below.

Another type of difference depends on whether the organisms can or cannot [alter](#) their information-processing capabilities over time (as in many types of learning), and whether and how they use information gained from the environment in order to do that, e.g. learning to perceive, learning to plan, learning to debug plans, learning to acquire and use geographical information (about an extended environment), learning to deal with conspecifics (including mates and offspring), potential prey, predators, collaborators, etc.

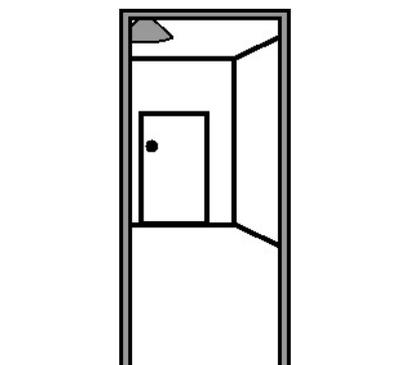
-----  
[Back to contents](#)  
-----

## From matter-manipulation to information-manipulation

Evolved versions of species with those capabilities might also develop abilities to acquire meta-information: not just information about locations, objects, motions, etc. but information about sources of information, about locations from which information is available, about new ways of transforming, combining and using information, and in some species abilities to pass information to conspecifics or to acquire information from them.

An example would be information about how to alter your own location so as to get new information, or more information, about part of the environment.

**Figure: Room view**



Which way could you move to see more of what's to the left of the visible door at the far end of the partly visible room?

Information about which movements will produce new information of desired kinds could be acquired experimentally -- by trying various kinds of motion.

However, an organism that knows that visual information travels in straight lines could **work out** which way to move to acquire more information about a particular currently invisible location, as in the door example above.

Compare that with situations in which instead of the perceiver moving to a new location, the object perceived can be moved or rotated.

Gaining information about how various movements produce changes in available visual information could be done by performing many experiments.

However, the information that visible information travels in straight lines could enable a perceiver with an understanding of spatial structures and relationships to **work out** how to move to acquire more information -- without having to use trial and error.

What sorts of evolutionary changes enabled our ancestors (and other ancient species)

- to use trial and error to acquire meta-information about how to get information (i.e. empirically)?
- to acquire meta-information about how to get information so that which way to move in a novel situation could be deduced, without empirical learning?

Other kinds of meta-information can be about how it is impossible to gain more information -- e.g. *rotating on the spot will give you more information about the room you are in, but not more information about the contents of a room seen through a distant doorway.* (Why not?)

The above examples provide a tiny subset of examples of possible evolutionary transitions in kinds of information-processing capabilities, including possible transitions that could have been precursors of abilities to make discoveries of the sorts assembled in Euclid's *Elements*.

Other important kinds of intelligence involving information about and reasoning about spatial structures and locations include information about extended terrain, with different resources, dangers and difficulties collected and assembled over time to form useful information about locations of resources and dangers and routes between them.

That could have co-evolved with spatial planning capabilities.

Different species have different subsets of such capabilities, including some acquired genetically (e.g. in migrating birds) others acquired by exploring the local environment.

### **Meta-reasoning**

Some organisms also have meta-reasoning capabilities required for acquiring and using information about information. An example is acquiring information about which way to move to acquire additional information, such as information about a partly visible surface.

There are likely to be many more types of spatial intelligence and many more explanatory mechanisms than we have so far thought of or been able to give to AI products.

Trying to learn about possible evolutionary trajectories may provide useful information about how such capabilities can vary, and whether they can be improved by learning or training (as opposed to re-designing).

By finding out what sorts of information processing capabilities our distant (non-human, and pre-human) ancestors had we may learn about previously unnoticed sub-competences that are still important.

By looking at intelligence in different evolutionary lineages (e.g. bird intelligence, octopus intelligence, slime mould intelligence) we may find clues about how different biological mechanisms (including different brain designs) can provide related competences -- helping us understand different design and implementation options and trade-offs for AI systems.

### **Mathematical competences**

In particular, I think we'll find evidence relating to the evolution of mathematical competences that help us find new (better) answers to old philosophical questions about the nature of mathematics, and the connection with abilities to acquire or derive and use information about what is and is not possible.

Knowledge of Euclidean geometry (and topology) can be understood as a product of particular types of evolutionary developments meeting particular challenges posed by spatial environments and their occupants.

But they were not arbitrary products of evolution: they relate to and are driven by deep mathematical properties of the environment, and the physical components of organisms, making some kinds of designs and behaviours possible and others impossible.

I have argued elsewhere that evolution makes use of mathematical features of our space-time world and its occupants in many of its designs, e.g. designs for control systems for organisms that change their size, shape, strength and sensory powers during individual development.

Moreover, because of the different instantiations of those general features in different environments, and at different stages of individual development, evolution produced genetic mechanisms that do not specify [detailed](#) designs for information processing within a species.

Instead, for some species, it produced designs that support [layered](#) forms of development, where different layers build partly on the need to meet more sophisticated information processing requirements, but also on information and abilities the individual has already acquired by interacting with the environment.

So different developmental stages may use different kinds and types of information about the environment and about previous developments in the individual, as sketched in [Figure Evo-Devo](#) below.

-----  
[Back to contents](#)  
-----

## EVO-DEVO Issues

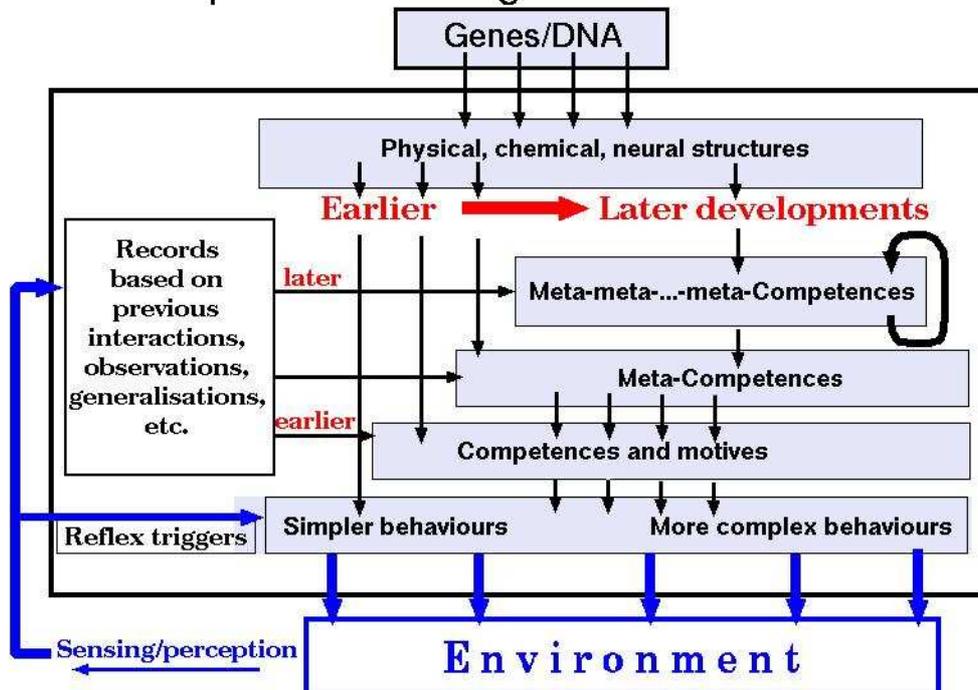
One of the important biological uses of information is in (partial) control of development by genetic information.

However, this control is often misconstrued as direct causation, ignoring the ways in which genetic information and environmental information are combined and coordinated -- but *not* like concurrently active forces of different strengths.

In 2005 and 2007 Jackie Chappell and I presented some ideas about this [[Chappell Sloman 2007](#)] which can be loosely summarised in this diagram:

**Figure Evo-Devo**

### Multiple routes from genome to behaviours



(This can be seen as a generalisation of Waddington's 'Epigenetic Landscape' -- for individuals that re-design the landscape during their own development.)

This diagram and the related published papers do not yet incorporate the additional varieties of information-flow from conspecifics to the individual or *vice versa* involving explicit teaching and learning, co-discovery, and cultural changes.

I think Jean Piaget understood much of this which is why he described the topic of his research as "Genetic epistemology" rather than developmental psychology.

Compare the work of Annette Karmiloff-Smith [[1992](#)].

See also, this discussion of Toddler Theorems:  
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>

## Implications for evaluation of theories of intelligence

We need a good theory not about a particular type of intelligent individual but a theory that explains how many varieties of natural and artificial intelligence are possible, including the intelligence of nest building birds, squirrels, elephants, human toddlers, and possible future robots, and future evolved species.

I have provided only a **subset** of ways in which types of intelligence in organisms and machines can vary. A good theory of intelligence would have to go into a lot more detail, and cover a lot more variety.

The **tests** required for those theories would be very different, e.g. testing whether a theory adequately explains intelligence involved in physical manipulations of various sorts vs testing whether a theory adequately explains how cognitive development happens.

I suggest that ALL so-called tests for intelligence that do not explicitly contribute towards the larger goal of testing a **theory** of intelligences of many types is unlike to contribute anything of lasting value to science or philosophy -- though it may serve some limited engineering goal (e.g. testing a virtual museum guide).

A generic theory would be related to particular varieties of intelligence somewhat as a general theory of chemistry relates to the chemical properties and capabilities of a huge variety of types of molecule and the reactions they can be involved in.

In contrast a particular test for intelligence would be as limited as a test for a particular type of chemical: hardly a great contribution to chemistry.

Tests for a general theory of intelligence would have to cover a wide variety of types of mind.

That would include types of mind that can discover and use mathematical features of the environment without being aware of what they are doing (e.g. children using topological features of the environment or grammatical features of their language), while other types of mind (e.g. some older humans) can notice and reflect on such mathematical features, e.g. making discoveries in geometry, topology and other branches of mathematics.

Biological evolution seems to have discovered a need for different levels and types of meta-cognition with different uses -- many of which have gone unnoticed by AI researchers, most psychologists and neuroscientists.

The problems include how to characterise and how to explain the mathematical abilities apparently shared between our ancestors and other intelligent animals, but not yet available in robots or AI theorem provers. In particular, we cannot yet explain or model the evolutionary processes that produced the capabilities used by the ancestors of Euclid, Archimedes and other ancient mathematicians (who knew nothing of modern logic, algebra, and formal proofs).

These questions are closely related to questions Immanuel Kant asked about mathematical knowledge around 1781 (in his *Critique of Pure Reason*).

Related but different capabilities lead to production of what we call "works of art", including music, dances, paintings, sculptures, etc. But the intelligence involved in such creative processes is different from the intelligence required to understand and enjoy the products.

There are now many AI programs that can produce musical or pictorial works of art, but as far as I know they cannot enjoy or appreciate their products.

What is missing? Are different mechanisms required for appreciating different types of art (or humour)?

This sort of motivation for research in AI differs from an interest in AI as [engineering](#), with the goal of producing [useful](#) intelligent systems, including machines that interact with and help humans. A recent article by Susan Epstein (Wanted Collaborative Intelligence, AIJ, 2015) illustrates some of the differences in requirements:

<http://www.cs.hunter.cuny.edu/~epstein/papers/AIJ%202015.pdf>

-----  
[Back to contents](#)  
-----

## Evolution as a blind mathematician

Evolution makes mathematical discoveries and uses them in many ways, including design of control systems, parametrised design for developing individuals, and many more.

Moreover its products express implicit theorems about what is mathematically possible.

The evolutionary and developmental trajectories are implicit proofs of those possibilities.

But some of evolution's **products** also make and use mathematical discoveries, including various mechanisms used for development and learning in individuals (some indicated in [Figure Evo-Devo](#), above).

Only humans seem to have the ability to notice that they are doing that, organise the discoveries, communicate them to others, challenge the claimed discoveries, and drive continual further developments of mathematics.

REF separate documents.

-----  
[Back to contents](#)  
-----

## Inspiration from Alan Turing

**If Turing had lived longer, how might he have investigated what AI and Philosophy can learn from evolved information processing systems?**

The proposed answer is presented in the form of conjectures about what he might have done if he had lived several more decades after publishing his paper on morphogenesis [Turing,1952].

The Meta-Morphogenesis (M-M) project is based on that conjectured answer. It is an attempt to identify forms and mechanisms of information processing produced by evolution since our planet formed, including forms of computation possibly relevant to deficiencies in current AI: for example deficiencies in spatial reasoning, and developmental and architectural deficiencies.

Many researchers, in many disciplines, have investigated changes produced by biological evolution, including: changes in modes of reproduction, changes in genetic material, changes in morphology (physical forms) of organisms, changes in sensory/perceptual mechanisms, changes of behaviour, changes of environment, changes in cultures, changes in ecosystems, and changes in modes of communication.

If Turing had lived several decades longer it seems likely, extrapolating from his previous interests and publications, that he would have done ground-breaking work on evolutionary changes in [types](#) of information, [uses](#) of information, and information-processing [mechanisms \(and architectures\)](#), from the very simplest life forms (or pre-life chemical structures) to present day organisms, and beyond, as conjectured here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

[Back to contents](#)

### **Evolution of mathematical abilities**

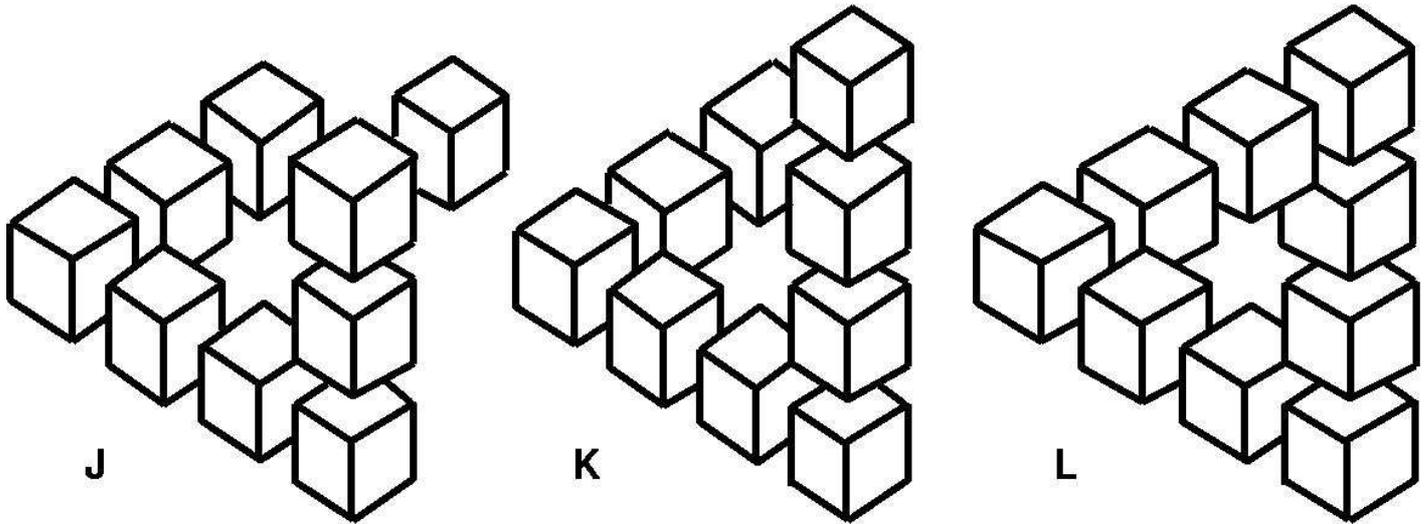
A subsidiary conjecture is that this sort of research, especially research leading to previously unnoticed transitions in forms of biological information processing, can help to shed light on some of the deep limitations of current AI.

The limitations include failures to identify or model important aspects of animal visual perception, and the forms of mathematical discovery and reasoning used by Euclid, Archimedes and their predecessors, and inability to model many of the competences and changes in competences in pre-verbal human children.

Important examples include abilities to see and reason not only about what currently exists in the environment but also about what changes are [possible](#), what changes are [impossible](#), and how some of the possible changes would alter what is and is not possible.

These abilities are shared across several species, and seem to be closely related to unexplained human abilities to make discoveries in geometry and topology of the sorts reported in Euclid's *Elements*.

**Figure: Blocks**



Yet more possible configurations of 9 blocks. Are they all really possible? See text for discussion.

(Inspired by [Reutersvard's 1934](#) drawing.)

An extended discussion and analysis can be found in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html#impossible>

-----

Such information about what is and is not possible, or necessarily the case, has nothing to do with statistical or probabilistic reasoning -- on which much (but not all) current AI is based.

Moreover, although there is a vast amount of research on brain mechanisms for statistical learning and reasoning, as far as I know nobody has any idea how brains represent necessity and impossibility.

It seems unlikely that they use modal logics with possible world semantics! I have discussed alternatives elsewhere including [\[Sloman 1962\]](#).

(The last two books by the major 20th Century developmental psychologist, Jean Piaget were concerned with how children develop in their ability to discover and reason about possibility and necessity, but I don't think he had deep ideas about explanatory mechanisms, though he discovered many fascinating things about children.)

Learning more about intermediate forms of those abilities in simpler organisms may draw attention to previously unnoticed aspects of brain function and gaps in AI mechanisms.

[Back to contents](#)

### **The need for construction kits**

Evolution produces many lineages of forms of life that assemble their own building blocks and information processing mechanisms (including control mechanisms) from resources available during individual development. For this it needs construction kits.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>

Products of evolution also make use of construction kits some of which are based on resources in the environment (e.g. solid surfaces, other organisms or products of other organisms) while others are assembled internally, e.g. the construction kits used to build the physical/chemical structures that enable many types of plant to grow upward and support themselves, from grasses to giant redwood trees.

More interesting for us are evolved construction kits for building information processing mechanisms.

I suspect there may be important special cases that have not yet been discovered by scientists or re-invented by engineers, including possibly some required for spatial reasoning such as mathematical discoveries in geometry and topology.

A partly developed theory of construction kits for evolution is presented here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>

[Back to contents](#)

### **Relevance to unsolved problems in other disciplines**

These topics are closely related not only to old problems in philosophy (e.g. about the nature of mind, language, and mathematical knowledge), but also to unsolved problems in developmental psychology and brain science. Connections will be made with the work of Immanuel Kant and Jean Piaget, among others.

Most AI researchers work on narrowly-focused (usually practical) problems and ignore these topics, without noticing what they are omitting. Two of the founders of AI, John McCarthy and Marvin Minsky, whose inspiring publications I read when I was a young philosopher learning about AI in the late 1960s, were not only interested in AI as engineering, but (like Alan Turing) also interested in AI as science and philosophy, as indicated by their publications and the willingness to take part in a session on AI and Philosophy at IJCAI in 1995 (see [above](#)).

### **The need for a theory of (non-Shannon) information and its role in the universe**

Discussed [below](#).

---

This was proposed as a half-day tutorial, but was accepted as a quarter-day (1 hour 45 Mins) tutorial.

Topics will be selected in part on the basis of responses from the audience.

---

[Back to contents](#)

-----

## **What Is Information?**

The universe contains Matter, Energy and Information.

Many important scientific concepts cannot be defined [explicitly](#), but are [implicitly](#) defined by their roles in theories.

The words "Matter", "Energy" and "Information", in the sense used here, cannot be explicitly defined but are **implicitly** defined by the theories that use the words, and the ways in which the theories are applied and tested.

Many researchers write as if the concept of information is essentially concerned with messages transmitted from a sender to a receiver, and stored before, during and after transmission. But those are secondary aspects of information. Studying them without studying uses of information will lead to distorted theories of information.

Any survey of biological information processing should be primarily concerned with the uses of information and how those uses differ during evolution and development, and between different species, and the resulting changing requirements for mechanisms involved in acquisition, manipulation, transformation, storage, analysis, combination, and above all **use** of different sorts of information. Investigating those changes over evolutionary time-scales is challenging because most of the evidence available must be indirect. Information and uses of information are not components of fossil or archeological records, for example, though they often provide clues.

Don't assume that the concepts we have now will prove adequate in the long run. Compare what happened to concepts of *force*, *weight* and *mass* between Newton and Einstein. I think our concepts and theories relating to information are still VERY primitive. (Though less primitive than Aristotle's theories of physics.)

Moreover research communities are dreadfully fragmented, using different systems of concepts with superficial overlaps at a verbal level.

In particular, many researchers assume that the only relevant concept of "information" is that presented by Claude Shannon in his 1948 two-part paper

A mathematical theory of communication,  
*Bell System Technical Journal*,

Unfortunately, because he was working for the Bell Telephone company Shannon's work was concerned with solutions to engineering problems related to storage and transmission of encodings of information in various forms. So he was not primarily concerned with the kind of information **content** that could be used, tested for truth or falsity or consistency with evidence or beliefs, or obeyed (e.g. instructions, or information about what to do). Shannon understood the difference but many of his admirers did not and were therefore seriously confused by his use of the term "information" forgetting that there was a much older use of the word, found for example in the novels of Jane Austen a century earlier, as shown in [[Austen\(Information\)](#)].

Storage and transmission are of interest only because **information can be used**: the primary role of information is to be used, not stored or transmitted. I think Shannon understood this, but his unfortunate terminology confused a whole generation of scientists and engineers.

For the M-M project, although we cannot offer an explicit definition of the key notion of "information" (any more than we can explicitly define "energy" or "matter") we can say much about the roles of information in living things.

The single most important point, as noted above, is that organisms USE information. The use or potential use is what makes storage and transmission of information important.

The use can take various forms, e.g. obeying an instruction, which involves performing an action corresponding to the content of the information (e.g. "If it is raining put on a coat"), using the information as a basis for choosing between possible actions (e.g. using information that a path is blocked as a basis for changing one's route). These two examples do not exhaust the types of use of information!

It is therefore a mistake to DEFINE information as something transmitted from a sender to a receiver. The roles of transmission, storage, retrieval, etc. are all secondary to the uses of information.

Specifying the variety of types of use of information, the variety of types of information, the mechanisms involved in use, the ways information can be acquired, and other such details will provide a "theory" of information that implicitly defines "information", as suggested in [Sloman\[2011 inf\]](#).

But the implicit definition can change and expand as our understanding of information, its uses, its forms, what can be done with it, etc. grows.

No deep concept in science can be fixed permanently by a prior definition. (E.g. see [Cohen 1962](#).) Compare the history of notions of force, mass, velocity, etc. in physics.

-----  
[Back to contents](#)  
-----

## **The fundamental use of information is for control**

Although there are many different uses of information, e.g. predicting, explaining, answering questions, some uses do not occur in the earliest life forms (e.g. explaining). So we need to understand the ways in which evolution produces new uses of information, and what the implications are, e.g. for sending, receiving, and storing information, while always remembering that they are secondary to uses. There would be no point sending, receiving and storing information if there was no use for it. Transmitting information and receiving information are among the processes that require explanatory mechanisms. But they are secondary to the mechanisms that use information: if there were no uses, there would be no need for transmission or storage.

At its most basic, information is used for control, i.e. to select, initiate, modify or terminate actions: internal or internal actions. This can take very simple forms, with "direct triggering" of a response by a recognition process, and increasingly complex forms as organisms get more complex, using more varieties of information more varied and complex purposes. An intermediate case is recognition of a need which indirectly produces action to meet the need as a result of a planning process. As more varied needs and collections of information become available, the mechanisms for acquiring, transforming, storing, retrieving, combining, and using information get more complex. Some of the mechanisms are concerned with the existence of multiple needs, opportunities and preferences, some of which can produce internal conflicts that can be detected by meta-cognitive (or meta-management) mechanisms. These in turn provide a potential basis for yet more complex multi-layered information processing mechanisms, some of which are outlined in the framework of the Birmingham CogAff (Cognition and Affect) project:

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

So any definition of information in terms of sending and receiving is misguided. Shannon unintentionally confused many people by choosing the label "information" for what he was writing about. (He was not confused).

Moreover information contents do not primarily vary in terms of how much of something they contain, but what they contain which can vary in structure, as sentences, maps, graphs, and other information representations do. So numerical *measures* of information may be relevant to requirements for storage or processing or transmission, but that does not make information measures relevant to content.

Moreover the standard measures can depend on forms of encoding, and the same information content can often be encoded in different ways, with different measures (e.g. size of bit string) but no difference in content.

There are very many forms of control, using different kinds of information, from different sources, for different purposes, employing different mechanisms. A very simple form of control is turning something on or off. More complex forms of control (e.g. in negative feedback controllers) use information to determine whether something should be increased or decreased (e.g. speed, pressure, temperature, curvature of a path, etc.).

One of the major features of biological evolution is production of increasingly sophisticated uses of information in increasingly complex mechanisms, for increasingly complex and varied purposes.

## **Evolutionary transitions from microbes to intelligent animals**

Between the simplest and most sophisticated organisms there are many intermediate forms with very different information processing requirements and capabilities.

The simplest organisms use information only for *online control*, e.g. in many biological forms of homeostasis.

The information in simple cases is constantly acquired used and discarded or over-written by new information.

More sophisticated organisms can acquire information for varied *offline* uses.

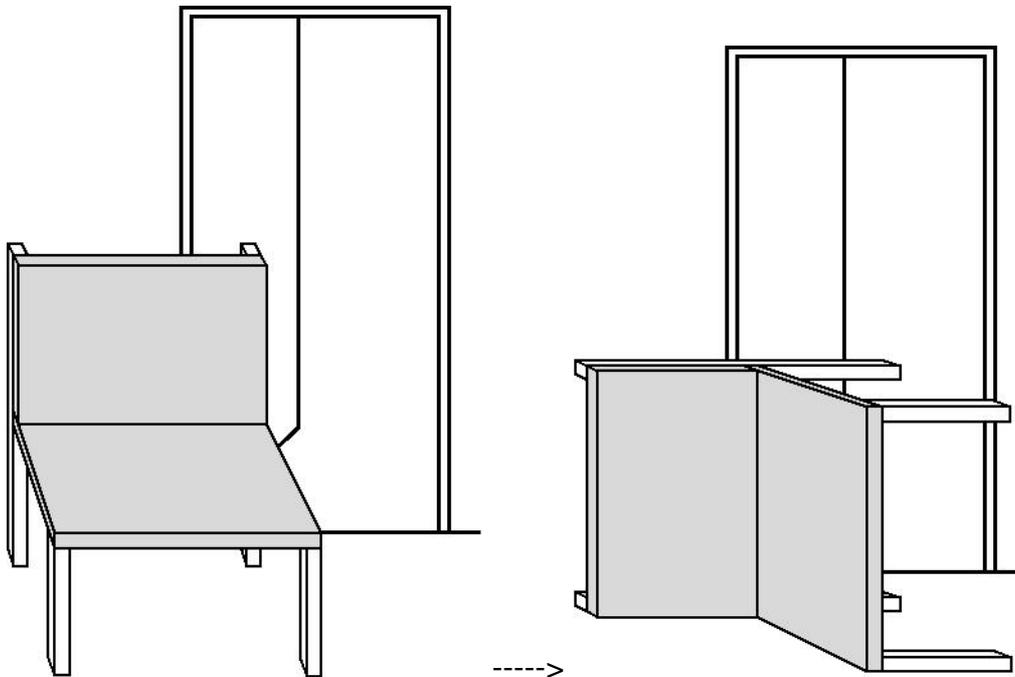
(There is much confusion about this: e.g. the spurious distinction between WHAT vs WHERE visual information in psychology and neuroscience.)

A particular type of off-line intelligence is the ability acquire and use information that is not merely about the **current** state of something (e.g. a part of the environment, or an adversary's goals), but also informs about possible changes of state, invariants across those changes and limitations of such changes.

An example would be acquiring information about possible routes by which an object could be transferred from one location to another, and constraints on those possibilities (e.g. the need for the object to have different spatial orientations at different stages of the motion) -- e.g. possible trajectories for a twig to be added to a part-built nest, or possible transitions of leaf to form a knot as part of a weaver-bird's nest.

Another example would be acquiring information about possible re-arrangements of a collection of objects to form a platform on which the maker could stand in order to reach something or to see over something: e.g. which movements of a pile of stones would form a platform on which you could stand in order to look over a high hedge.

Or information about how to re-orient a chair so that it can be moved through a doorway:



These are all examples of uses of intelligence in selecting and guiding processes in which physical matter is manipulated.

The uses of information, the information contents, the forms of representation or encoding, and the forms of control vary enormously across evolutionary time scales, across species, across the life-span of individual organisms, and between functions of the whole organism (e.g. perception and learning) and microscopic control processes deep within organisms, e.g. concerned with immune responses, digestion, temperature control, control of growth, reactions to infection or damage, varieties of metabolism, and many more.

Not all of these have been considered relevant to AI, but they become increasingly relevant as robots become more sophisticated and use more complex components.

Many of the research questions concern the variety of types of information about the environment and the variety of uses to which the information can be put, e.g. immediate decision making or variation of control parameters during action, vs finding out what is and is not possible in a situation, and how possibilities and impossibilities change if some possibilities are realised. This is obviously relevant to decades-old research on planning, plan execution, action control.

Less obviously it is relevant to the kinds of discovery our ancestors made about geometry and topology, and the discoveries young children and other species make, without being aware of what they are learning. I'll present examples in the tutorial, but some samples are available here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html>

## What don't we understand yet

It is not clear to me that the forms of representation so far investigated in AI, robotics, psychology, and neuroscience, are adequate for the tasks. I'll give examples related to the need to explain both examples of child development studied by Piaget and others and the need to explain evolution of mathematical abilities of Euclid etc.

This is also relevant to problems in philosophy of mathematics discussed by Immanuel Kant in his *Critique of Pure Reason* (1781).

It is relevant to problems of varieties of modality and how they are used, studied in natural language processing, theoretical linguistics, logic, semantics, ontology,... Does a child or squirrel discovering that something is possible, or impossible, need to think about possible worlds, or is there a deeper, more realistic explanation of how modal concepts work (possible, impossible, necessary, contingent)?

-----  
[Back to contents](#)  
-----

## Criticisms of some current fashions

Depending on time available and interests of participants I may be able to fit in some of these topics.

### **Rich internal languages must have evolved before languages for communication**

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>

### **The shallowness of some recent AI (or anti-AI) fashions: embodiment, enactivism, extended-mind, ...**

(Semi-spoof: the Chewing Test for Intelligence.)

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/chewing-test.html>

### **The problem of specifying the functions of vision in intelligent systems**

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision/plants/>

### **More on Natural and Artificial vision**

The over-emphasis of numerical and statistical techniques in AI instead of use of relational descriptions, partial orderings, and representations of possibilities and impossibilities.

Why the use of kinect instead of cameras may be depriving robots of useful sources of information about the environment, its affordances, changes of affordances, and their own actions.

Alternative (non-probabilistic) ways to deal with uncertainty.

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#718>

### **The need for a broader view of the functions of vision**

(Including alternative ways to do stereo vision)

How Julesz random-dot stereograms lead researchers to poor theories of stereo vision, ignoring rich monocular structure in many natural images.

The roles of biological vision in mathematical discovery (e.g. geometry, and topology)

Additional challenges for machine vision and its roles in cognition.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html>

### **Varieties of motivation**

(To be expanded)

Architecture-based motivation (ABM) vs Reward-based motivation (RBM) [[Sloman 2009](#)]

### **Architectures**

The need for a principled theory of varieties of information processing architectures for more or less intelligent animals and machines -- including a view of architectural changes across evolutionary time scales and architectural changes during individual development.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/fully-deliberative.html>

### **Virtual machines and consciousness**

My talk presupposes that understanding the roles of various kinds of virtual machinery can give us new answers to old problems about consciousness, in animals and machines, as discussed here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>

It will not be possible to discuss all the above topics in the time allocated for the tutorial. I may use audience reaction to tailor the selection of topics for presentation to the interests (or knowledge-gaps) in the audience. Anyone interested in following up sub-topics can explore web pages linked from this one and email me with comments and questions. There may also be time to continue talking after the formal close of the tutorial.

## **Sample References**

---

### **Note on "Bridging" workshop paper.**

Ideas relevant to the tutorial are in a paper for the workshop on [W1: Bridging the Gap between Human and Automated Reasoning](#), held on the preceding day, Saturday July 9th, 2016. Details at:

<http://ratiolog.uni-koblenz.de/bridging2016#Progr>

<http://ratiolog.uni-koblenz.de/bridging2016#Papers>

My paper for that workshop is directly relevant to this tutorial.

[Sloman(Bridging)]

A. Sloman, Natural Vision and Mathematics: Seeing Impossibilities

(About human abilities to make discoveries in geometry and topology, and related abilities in other intelligent animals -- abilities not yet available to AI reasoning systems.)

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-bridging-gap-2016.pdf>

---

[Austen(Information)]

Jane Austen's concept of information (As opposed to Claude Shannon's) (Online discussion note.)

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html>

[Beaudoin, 2014]

L.P. Beaudoin, 2014, *Cognitive Productivity: The Art and Science of Using Knowledge to Become Profoundly Effective*, Leanpub. <http://leanpub.com/cognitiveproductivity>

[Beaudoin, CogZest]

Luc Beaudoin (2015??), *The CogZest web site* <https://cogzest.com/>  
<https://cogzest.com/2015/07/a-tale-of-two-summer-conferences-isre-2015-and-cogsci-2015/>

[Chappell Sloman 2007]

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/cogaff/07.html#717>

[Cohen 1962]

L.J. Cohen, 1962, *The diversity of meaning*, Methuen & Co Ltd, London

[Crick, 1954]

F. H. C. Crick, 1954/2015 The structure of the hereditary material, in *Nobel Prizewinners who changed our world*  
Scientific American, Topix Media Lab, New York USA 1954/2015 pp. 6--15

[Ganti,Life]

Tibor Ganti, 2003. *The Principles of Life*,  
Eds. E. Szathmáry, & J. Griesemer, (Translation of the 1971 Hungarian edition), OUP, New York.  
See the very useful summary/review of this book by Gert Korthof:  
<http://wasdarwinwrong.com/korthof66.htm>

[Gibson 1979]

J.J. Gibson (1979) *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA

[Glasgow, Narayanan, Chandrasekaran; . 1995]

Glasgow, J., Narayanan, H. Chandrasekaran, B. (). (Eds.). (1995). *Diagrammatic reasoning: Computational and cognitive perspectives*. Cambridge, MA: MIT Press.

[ Jablonka Lamb 2005]

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 1781]

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

[Lakatos 1976]

Lakatos, I. (1976). *Proofs and Refutations*. Cambridge, UK: Cambridge University Press.

[Marr 1982]

Marr, D. (1982). *Vision*. San Francisco: W.H.Freeman.

[ McCarthy HayesMcCarthy Hayes 1969]

McCarthy, J. Hayes, P. (1969). Some philosophical problems from the standpoint of AI. In B. Meltzer & D. Michie (Eds.), *Machine Intelligence 4* (pp. 463-502). Edinburgh, Scotland: Edinburgh University Press.

[Piaget 1952]

Jean Piaget, 1952, *The Child's Conception of Number*, Routledge & Kegan Paul, London,

[Piaget 1981-1983]

Jean Piaget (1981,1983) *Possibility and Necessity* (1981/1983) *Vol 1. The role of possibility in cognitive development* (1981), *Vol 2. The role of necessity in cognitive development* (1983), (Tr. by Helga Feider from French in 1987)

[Rescorla 2015]

Michael Rescorla, (2015), The Computational Theory of Mind, in *The Stanford Encyclopedia of Philosophy* Ed. Edward N. Zalta, Winter 2015,

<http://plato.stanford.edu/archives/win2015/entries/computational-mind/>

[Senghas 2005]

Senghas, A. (2005). Language Emergence: Clues from a New Bedouin Sign Language. *Current Biology*, 15 (12), R463-R465.

[Sloman 1962]

Sloman, A. (1962). Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth (DPhil Thesis).

<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962>

[Sloman 1971]

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/62-80.html#1971-02>

(Reprinted in *Artificial Intelligence*, vol 2, 3-4, pp 209-225, 1971)

[Sloman 1978 1]

Sloman, A. (1978a). The computer revolution in philosophy. Hassocks, Sussex: Harvester Press (and Humanities Press). <http://www.cs.bham.ac.uk/research/cogaff/62-80.html#crp>

[Sloman 1978 2]

Sloman, A. (1978b). What About Their Internal Languages? Commentary on three articles by Premack, D., Woodruff, G., by Griffin, D.R., and by Savage-Rumbaugh, E.S., Rumbaugh, D.R., Boysen, S. in *BBS Journal* 1978, 1 (4). *Behavioral and Brain Sciences*, 1(4), 515.

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#713>

[Sloman 1979]

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/62-80.html#1979-01>

[Sloman 1982]

Sloman, A. (1982). Towards a grammar of emotions. *New Universities Quarterly*, 36(3), 230-238.

<http://www.cs.bham.ac.uk/research/cogaff/81-95.html#emot-gram>

[Sloman 1983]

Sloman, A. (1983). Image interpretation: The way ahead? In O. Braddick & A. Sleight. (Eds.), *Physical and Biological Processing of Images* (Proceedings of an international symposium organised by The Rank Prize Funds, London, 1982.) (pp. 380-401). Berlin: Springer-Verlag.

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#57>

[Sloman 1984]

Sloman, A. (1984). The structure of the space of possible minds. In S. Torrance (Ed.), *The mind and the machine: philosophical aspects of artificial intelligence*. Chichester: Ellis Horwood.

<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#704>

[Sloman 1985]

Sloman, A. (1985). What enables a machine to understand? In *Proc 9th IJCAI* (pp. 995-1001). Los Angeles.

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#4>

[Sloman 1986]

Sloman, A. (1986). What Are The Purposes Of Vision? (Presented at: Fyssen Foundation Vision Workshop Versailles France, March 1986, Organiser: M. Imbert)

<http://www.cs.bham.ac.uk/research/projects/cogaff/12.html#1207>

[Sloman 1989]

Sloman, A. (1989). On designing a visual system (towards a Gibsonian computational model of vision). *Journal of Experimental and Theoretical AI*, 1(4),

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#7>

[Sloman 1992]

Sloman, A. (1992). The emperor's real mind. *Artificial Intelligence*, 56, 355-396. (Review of Roger Penrose's *The Emperor's new Mind: Concerning Computers Minds and the Laws of Physics*)

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#12>

[Sloman 2006]

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

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/fully-deliberative.html>

[Sloman 2009]

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

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html>

[Sloman 1996]

A. Sloman, 1996, Beyond Turing Equivalence, in *Machines and Thought: The Legacy of Alan Turing (vol I)*, Eds. P.J.R. Millican and A. Clark, The Clarendon Press, Oxford, pp. 179--219,

<http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#1>

[Sloman 2009]

A. Sloman, 2009, Why the "hard" problem of consciousness is easy and the "easy" problem hard. (And how to make progress), Online tutorial presentation,

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#cons09>

University of Birmingham.

[Sloman, Vernon 2007]

Aaron Sloman and David Vernon, (2007). A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots, Contribution to euCognition wiki

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-requirements.html>

[Sloman 2010 1]

Sloman, A. (2010a, August). How Virtual Machinery Can Bridge the "Explanatory Gap", In *Natural and Artificial Systems*. In S. Doncieux & et al. (Eds.), *Proceedings SAB 2010*, LNAI 6226 (pp. 13-24). Heidelberg: Springer.

<http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#sab>

[Sloman 2010 2]

Sloman, A. (2010b). If Learning Maths Requires a Teacher, Where did the First Teachers Come From? In A. Pease, M. Guhe, & A. Smaill (Eds.), *Proc. Int. Symp. on Mathematical Practice and Cognition*, AISB 2010 Convention (pp. 30-39). De Montfort University, Leicester.

<http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1001>

[Gibson, 1979]

Gibson J.J. (1979) *The Ecological Approach to Visual Perception*. Houghton Mifflin, Boston, MA

[Huxley 1866/1872]

T. H. Huxley, *Lessons in Elementary Physiology*, MacMillan and Co, New York, 1866. 7th Edition 1872. <http://aleph0.clarku.edu/huxley/Book/PhysioL.html>

[Kant 1781]

Immanuel Kant (1781), *Critique of Pure Reason*, Translated (1929) by Norman Kemp Smith, London, Macmillan.

[Karmiloff-Smith 1992]

Annette Karmiloff-Smith, 1992, *Beyond Modularity: A Developmental Perspective on Cognitive Science*, MIT Press.

Partial review/discussion:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html>

[Minsky 1987]

Marvin L. Minsky, 1987, *The Society of Mind*, William Heinemann Ltd., London.

[Minsky 2006]

Marvin L. Minsky, 2006, *The Emotion Machine*, Pantheon, New York

[Sloman 1962]

Aaron Sloman (1962), *Knowing and Understanding: Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth*. Oxford University DPhil Thesis, 1962. Recently digitised (2016):

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-1962>

[Sloman et al 1991--]

Aaron Sloman, students and colleagues (1991--present) *The Cognition and Affect Project Web site*

<http://www.cs.bham.ac.uk/research/projects/cogaff/#overview>

[Sloman,2009]

Aaron Sloman, 2009, Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress, in *Creating Brain-like Intelligence*, Eds. B. Sendhoff, E. Koerner, O. Sporns, H. Ritter and K. Doya, Springer-Verlag, pp. 248--277, Berlin, <http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#912>

[Sloman-Toddler]

Aaron Sloman (2011-2016). *Meta-Morphogenesis and Toddler Theorems: Case Studies*, Online technical report (HTML and PDF), School of Computer Science, University of Birmingham. <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>

[Sloman (2011 Inf)]

A. Sloman (2011) 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, pp.393--438, <http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#905>

[Sloman(Vision)]

Sloman, A. (2011, Sep). What's vision for, and how does it work? From Marr (and earlier) to Gibson and Beyond. (Online tutorial presentation, also at <http://www.slideshare.net/asloman/>) <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk93>

[Sloman-Poly]

Aaron Sloman (2011-...), Family Resemblance vs. Polymorphism A comparison: (Wittgenstein's Family Resemblance Theory vs. Ryle's Polymorphism and Polymorphism in Computer Science/Mathematics.) Online discussion note, University of Birmingham (2011-2016). <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html>

[Sloman 2013]

Sloman, A. (2013). Virtual machinery and evolution of mind (part 3) Meta-morphogenesis: Evolution of information-processing machinery. In S. B. Cooper & J. van Leeuwen (Eds.), *Alan Turing - His Work and Impact* (p. 849-856). Amsterdam: Elsevier. <http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d>

[Sloman and Croucher 1981]

Sloman, A., & Croucher, M. (1981). Why robots will have emotions. In *Proc 7th int. joint conference on AI* (pp. 197-202). Vancouver: IJCAI. <http://www.cs.bham.ac.uk/research/cogaff/81-95.html#36>

[Turing,1952]

A. M. Turing (1952), The Chemical Basis Of Morphogenesis', in *Phil. Trans. R. Soc. London B* 237, 237, pp. 37--72, 1952

[SEP "Consciousness"]

Robert Van Gulick, Robert, "Consciousness", *The Stanford Encyclopedia of Philosophy* (Spring 2014 Edition), (Ed) Edward N. Zalta. Spring 2014 (substantive content change) <http://plato.stanford.edu/archives/spr2014/entries/consciousness>

[von Neumann 1958]

John von Neumann *The Computer and the Brain (Silliman Memorial Lectures)*, 1958 (Yale University Press, 2012, 3rd Edition, with Foreword by Ray Kurzweil),

---

## **ORIGINAL DRAFT PROPOSAL FOR HALF-DAY TUTORIAL (Later revised to fit quarter-day time slot.)**

### **1. Brief description of the tutorial**

The tutorial will consist of a highly interactive discussion of a proposed answer to the question: "What might Alan Turing have worked on if he had lived several more decades after publishing his 1952 paper on Morphogenesis?"

Proposed answer: He might have worked on something like the (Turing-inspired) *Meta-Morphogenesis* project, which attempts to identify forms and mechanisms of information processing produced by natural selection since our planet formed, perhaps including forms of computation related to major deficiencies in current AI.

**Presenter:** Aaron Sloman (<http://www.cs.bham.ac.uk/~axs>)  
[Biography below](#)

### **2. Longer overview of the tutorial.**

Turing's paper on "The chemical basis of morphogenesis" was published in 1952. Two years later he was dead. What might he have done if he had lived several decades longer?

An answer was proposed in [Sloman 2013] and later developed in [goo.gl/9eN8Ks](http://goo.gl/9eN8Ks) that seems to be consistent with the questions Turing had asked since childhood, namely: Turing might have worked on something like the Meta-Morphogenesis (M-M) project: a highly ambitious, very long term, multi-disciplinary attempt to identify and learn from forms of information-processing that evolved on Earth.

Between the (partly understood) original chemical information processing mechanisms in the simplest life (or proto-life) forms and the most advanced known brain mechanisms there may have been many intermediate forms of information processing, not yet recognized by human scientists. Perhaps some of the more important intermediaries still play an unrecognized role in brains of recently evolved highly intelligent organisms, including squirrels, crows, elephants, cetaceans and apes, all of whom seem to have capabilities still unmatched by AI systems.

How could all that have emerged from a lifeless planet?

What will it take to replicate the evolved functionality using human-designed machinery? Some researchers attempt to answer the question by assuming that the functions of individual neurons are comparable to the functions of individual transistors in a computer, and then use Moore's Law (already defunct!) to predict when computers will match the power of human brains.

But if chemistry-based computation is important in brains (as suggested by Turing in his *Mind* (1950) paper) then the required number of transistors could be several orders of magnitude larger and very much longer times will be required to replicate brain functions in human-made systems. John von Neumann recognized that possibility in his book *The Computer and the Brain* (written in 1956 for the Silliman Memorial Lectures while he was dying of cancer, and published in 1958). The

argument is summarised by Tuck Newport in his little book *Brains and Computers: Amino Acids versus Transistors*  
<https://www.amazon.com/Brains-Computers-Amino-versus-Transistors-ebook/dp/B00OQFN6LA/>

In 1944 Schrödinger pointed out (in *What is life?*) that in addition to the non-determinism and statistical nature of quantum phenomena, at molecular scales there are highly deterministic discrete mechanisms without which life as we know it would be impossible: part of the Fundamental Construction Kit (FCK) provided by physics, which, as he noted, had important information-processing capabilities (anticipating digital computers).

Commented extracts from *What is life?* are available here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/schrodinger-life.html>

A draft introductory (high level) survey of types of construction-kit used by evolution is here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/construction-kits.html>

We discuss information-processing where "information" is understood as in Jane Austen's novels rather than Shannon's sense:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html>

Evolution produced many more construction kits: Derived Construction Kits (DCKs). Different DCKs emerged in different evolutionary lineages.

Besides the **concrete** construction-kits built from physics and chemistry, evolution also produced and used **abstract** construction kits and **hybrid** Concrete+Abstract construction kits, including **Meta**-Construction kits for producing new construction kits during the development of individuals and communities (e.g. language construction kits, theory construction kits, ontology construction kits).

The tutorial will report on progress in developing these ideas, including ideas about the differences that can be made by new construction kits and new information-processing architectures that grow themselves -- suggesting a (much?) richer variety of forms of computation than the Church-Turing thesis (apparently not believed by Turing) suggests.

Examples will be presented where current AI does not seem to be close to producing adequate models or explanations, e.g. how our ancestors made the discoveries leading up to Euclid's *Elements*, or how pre-verbal children create languages, make topological and other discoveries, and develop meta-cognitive capabilities, and how animal vision works. The audience will be invited to contribute examples, relevant ideas, criticisms, and suggestions for speeding up progress. Further information about the M-M project, with links to many sub-topics and other disciplines can be found at <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

This project has deep connections with philosophy, expanding ideas proposed in 1978 in *The Computer Revolution in Philosophy*: <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/>

### 3. Outline of the tutorial.

This was a proposal for a half-day tutorial. A subset of this material will be selected for the shorter tutorial. (It may be possible for interested participants to continue discussion after the tutorial.)

Topics on the day will be selected from the contents of **2.** above, plus a subset of the following. (See also the M-M project web site.)

I believe that AI as *Science* has lost its way in the last two or three decades (partly as a result of impatience for results) and is not making progress with some of its hardest problems, despite spectacular narrowly focused successes, and many engineering successes.

Some examples of what's still missing from AI (but not impossible, I hope):

A: modelling/replicating animal/human perception of complex irregular structured scenes in motion, like the changing view of a garden full of different flowers, bushes, shrubs and trees, as you walk through it during an irregular breeze, or the changing views a nest-building bird (e.g. crow) gets as it returns to its nest, from different directions. Example videos and some challenges are here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vision/plants>

B: explaining the forms of discovery and reasoning that led to all the mathematical knowledge compiled in Euclid's *Elements* about 2,500 years ago before modern logic, formal systems, or the arithmetisation of geometry.

C: the ability of Nicaraguan deaf children to start learning a sign language and then go on to create a much richer one themselves, because the teacher's linguistic competence did not suffice for their needs.

See <https://www.youtube.com/watch?v=pjtioIFuNf8>

D: the forms of exploration and apparent types of discovery and types of self-extension demonstrated in young children even before they have the linguistic competences to express what they have learnt, or how they are using the knowledge.

E: many forms of intelligence in non-human animals, including squirrels, nest-building birds, hunting mammals, elephants, monkeys and apes.

F: Abilities to *enjoy* and to *create* various art forms including music, poetry, dance, painting, and stories, and to dislike others.

G: Abilities to develop new types of learning, individual and cooperative.

Being retired, I have been working on this almost full time since early 2012, and there is a large messy and still growing web site presenting the problems addressed in the Meta-Morphogenesis project, and various fragments of progress: <http://goo.gl/9eN8Ks>

One of the important themes of the project concerns evolution's use of construction kits. It started with a "Fundamental Construction Kit" (FCK) provided by physics and chemistry, with some very important features noted by Schrödinger in his little book "What is life" (mentioned above). Evolution built on top of the FCK many branching layers of "Derived Construction Kits" (DCKs) of many kinds, some of which are mainly concerned with building physical structures and mechanisms found in organisms while others are primarily concerned with building increasingly sophisticated types of information-processing mechanism.

It may take decades (or perhaps centuries?) to replicate enough of these mechanisms in artificial systems to account for the many aspects of animal intelligence that current AI does not seem close to achieving, including modelling Euclid's precursors who made mathematical discoveries before there were mathematics teachers.

All of this relates closely to questions in biology, psychology, evolution, linguistics, education and especially philosophy, including philosophy of mind, philosophy of science, epistemology, philosophy of language and philosophy of mathematics, topics that I first started trying to relate to AI in my 1978 book, *The Computer Revolution in Philosophy: Philosophy science and models of mind*, now freely available online here (partly revised): <http://goo.gl/AJLDih>

#### **4. Potential target audience.**

This should be relevant to anyone interested in AI as Science, ways in which AI as engineering has not yet replicated natural intelligence, the potential for AI to learn from other disciplines, including Biology, and the potential of other disciplines to learn from AI.

In particular this will address deep philosophical problems about AI and the nature of mind, including the variety of types of natural mind.

#### **5. Why the tutorial topic would be of interest to a substantial part of the IJCAI audience, and which of the above objectives are best served by the tutorial.**

This is answered in section 4 and earlier references to unsolved problems in AI.

### **6. INFORMATION ABOUT THE PRESENTER**

#### **A brief resume**

(Part of the original tutorial proposal.)

CV with partial publications list available online

<http://www.cs.bham.ac.uk/~axs/shortcv.html>

#### **6.a. Name and address**

Aaron Sloman  
School of Computer Science,  
The University of Birmingham  
Edgbaston, Birmingham, B15 2TT, UK  
<http://www.cs.bham.ac.uk/~axs>

#### **6.b. Background and publications/presentations**

I have been working on the overlap between AI and Philosophy since 1969 and presented a paper pointing out gaps in Logicist AI at the 2nd IJCAI in 1971, reprinted in AIJ the same year.

*Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence*

I helped to develop major AI/Cognitive Science teaching and research centres first at Sussex University (1972-1991) then at Birmingham University. That included major contributions to the Poplog System, and the SimAgent toolkit. <https://en.wikipedia.org/wiki/Poplog>  
<http://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html>

I have organised a number of workshops and Tutorials at IJCAI, AAI and other AI conferences.

I gave a half day tutorial at AAI in San Francisco in 2011, before the Meta-Morphogenesis project had begun:

<http://www.cs.bham.ac.uk/research/projects/cogaff/aaaitutorial/>

Philosophy as AI and AI as Philosophy

### 6.c. Example of work in the area

See [Sloman 2013] and the Meta-Morphogenesis web site: <http://goo.gl/9eN8Ks>

---

### Biography of presenter (requested for tutorial proposals):

Born 1936, QueQue Southern Rhodesia (now Zimbabwe). BSc mathematics and physics Cape Town (1956). Rhodes Scholar, Oxford 1957-60. Switched to Philosophy. DPhil on Kant's philosophy of mathematics 1962. While a lecturer at Sussex University, learnt about AI from [Max Clowes](#) in 1969, and decided that doing AI was the best way to answer certain philosophical problems, e.g. about minds, knowledge, language, mathematics and science. (Later that included [emotions](#), and other forms of affect.)

Spent 1972-3 with AI researchers in Edinburgh. At Sussex helped to start undergraduate courses in AI and later the School of Cognitive and Computing Sciences. Published "The Computer Revolution in Philosophy" in 1978. ([goo.gl/AJLDih](http://goo.gl/AJLDih))

Helped with development of the Pop-11 programming language, the Poplog development environment, the SimAgent toolkit, and much AI teaching material.

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

GEC-funded research professor 1984-6. Moved to Birmingham University, UK, in 1991.

Officially retired 2002 but continues doing research full time as Honorary Professor of AI and Cognitive science.

1991: Elected fellow of AAI (Association for the Advancement of AI)

1997: Elected honorary life fellow of AISB

1999: Elected fellow of ECCAI European Coordinating Committee on AI

2006: Honorary DSc, awarded July 2006, Sussex University

Now working mainly on the Meta-Morphogenesis project: <http://goo.gl/9eN8Ks>

Recently digitised 1962 DPhil Thesis: "Knowing and Understanding"

<http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-1962>

---

Saturday 7/9 - Monday 7/11 Workshops and Tutorials

<http://ijcai-16.org/venue.html>

[All accepted workshops](#)

[All accepted tutorials](#)