Why is it so hard to make human-like AI (robot) mathematicians? Especially Euclidean geometers.

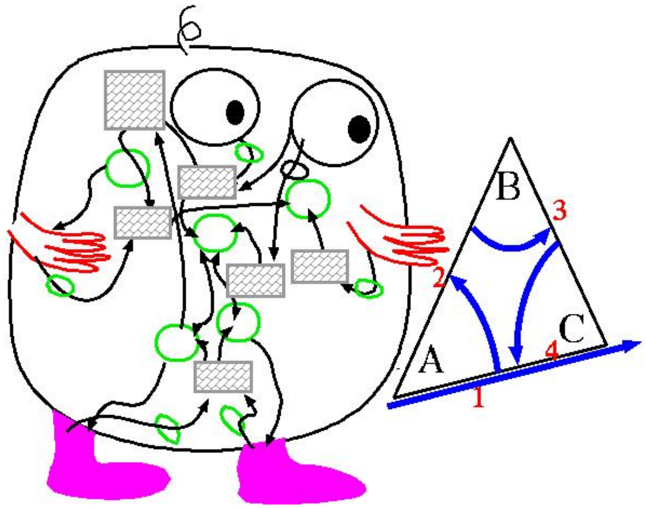

Closely related to interview and tutorial at AGI 2012, Oxford December 2012

Video of AGI Interview http://www.youtube.com/watch?v=iuH8dC7Snno

Video of AGI Tutorial http://www.youtube.com/watch?v=BNul52kFI74

Aaron Sloman

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

School of Computer Science, University of Birmingham.

This will later be added to my collection of online presentations

http://www.cs.bham.ac.uk/research/projects/cogaff/talks/
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk108

(Thanks to Marcin Milkowski for help in preparation.)

(X) Abstract As advertised

I originally got involved in AI many years ago, not to build new useful machines, nor to build working models to test theories in psychology or neuroscience, but with the aim of addressing philosophical disagreements between Hume and Kant about mathematical knowledge, in particular Kant's claim that mathematical knowledge is both non-empirical (apriori, but not innate) and non-trivial (synthetic, not analytic) and also concerns necessary (non-contingent) truths.

I thought a "baby robot" with innate but extendable competences could explore and learn about its environment in a manner similar to many animals, and learn the sorts of things that might have led ancient humans to discover Euclidean geometry.

The details of the mechanisms and how they relate to claims by Hume, Kant, and other philosophers of mathematics, could help us expand the space of philosophical theories in a deep new way.

Decades later, despite staggering advances in automated theorem proving concerned with logic, algebra, arithmetic, properties of computer programs, and other topics, computers still lack human abilities to think geometrically, despite advances in graphical systems used in game engines and scientific and engineering simulations. (What those do can't be done by human brains.)

I'll offer a diagnosis of the problem and suggest a way to make progress, illuminating some unobvious achievements of biological evolution.

The Meta-Morphogenesis Project

(Partly inspired by Turing's work on morphogenesis)

http://tinyurl.com/CogMisc/misc/meta-morphogenesis.html

How could our minds and the rest of life

have come from a cloud of dust?

Including: Microbes, mice, monkeys, mathematics, music, marmite (along with murder, mrsa, religious bigotry, and other nastiness).

Talk Abstract here
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bath-brlsi-talk.html

---

VIDEOS:

Interview with Adam Ford

http://www.youtube.com/watch?v=iuH8dC7Snno

Tutorial at AGI 2012

http://www.youtube.com/watch?v=BNul52kFI74

(PERMANENTLY) DRAFT DISCUSSION PAPERS:

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html

Main points 1a

AI, Robotics, Cognitive Science, Neuroscience, Psychology, Education, Evolutionary Biology, Ethology, and Philosophy

are all missing something that is important for all of them.

(Though various fragments have been noticed separately by different researchers)

Main points 1b

AI, Robotics, Cognitive Science, Neuroscience, Psychology, Education, Evolutionary Biology, Ethology, and Philosophy

are all missing something that is important for all of them.

Namely:

The nature of human mathematical competences thousands of years ago that led to the development of Euclidean Geometry

as presented in Euclid's **Elements**
(arguably the single most important book ever written on this planet).

Main points 1c

AI, Robotics, Cognitive Science, Neuroscience, Psychology, Education, Evolutionary Biology, Ethology, and Philosophy

are all missing something that is important for all of them.

Namely:

The nature of human mathematical competences thousands of years ago that led to the development of Euclidean Geometry ...

as presented in Euclid's **Elements**,
(arguably the single most important book ever written on this planet).

Which I suspect grew out of (proto-)mathematical spatial reasoning competences that evolved partly independently in different lineages.

Main points 1d

AI, Robotics, Cognitive Science, Neuroscience, Psychology, Education, Evolutionary Biology, Ethology, and Philosophy

are all missing something that is important for all of them.

Namely:

The nature of human mathematical competences thousands of years ago that led to the development of Euclidean Geometry

as presented in Euclid's **Elements**.
(arguably the single most important book ever written on this planet).

Which I suspect grew out of (proto-)mathematical spatial reasoning competences that evolved partly independently in different lineages.

Including: several bird species, elephants, cetaceans, apes, and (perhaps via apes?) humans.

Main points 1e

AI, Robotics, Cognitive Science, Neuroscience, Psychology, Education, Evolutionary Biology, Ethology, and Philosophy

are all missing something that is important for all of them.

Namely:

The nature of human mathematical competences thousands of years ago that led to the development of Euclidean Geometry

as presented in Euclid's **Elements**.
(arguably the single most important book ever written on this planet).

Which I suspect grew out of (proto-)mathematical spatial reasoning competences that evolved partly independently in different lineages.

Including: several bird species, elephants, cetaceans, apes, and (perhaps via apes?) humans.

Distractions from these hard problems include various fashions that inflate minor aspects of biological intelligence.

Main points 1f

AI, Robotics, Cognitive Science, Neuroscience, Psychology, Education, Evolutionary Biology, Ethology, and Philosophy

are all missing something that is important for all of them.

Namely:

The nature of human mathematical competences thousands of years ago that led to the development of Euclidean Geometry

as presented in Euclid's **Elements**.
(arguably the single most important book ever written on this planet).

Which I suspect grew out of (proto-)mathematical spatial reasoning competences that evolved partly independently in different lineages.

Including: several bird species, elephants, cetaceans, apes, and (perhaps via apes?) humans.

Distractions from these hard problems include various fashions that inflate minor aspects of biological intelligence.

embodiment, enactivism, extended mind, probabilistic learning, utility functions,

and the use of cartesian coordinate representations of spatial structures and processes.

Examples of Euclidean reasoning

This is just a tiny subset of examples: thousands more could be given.

Why do internal angles of a triangle add up to half a rotation?

How many people can see how to prove that?

Examples of Euclidean reasoning 1

This is just a tiny subset of examples: thousands more could be given.

Why do internal angles of a triangle add up to half a rotation?

A proof due to Mary Pardoe (when teaching maths about 40 years ago)

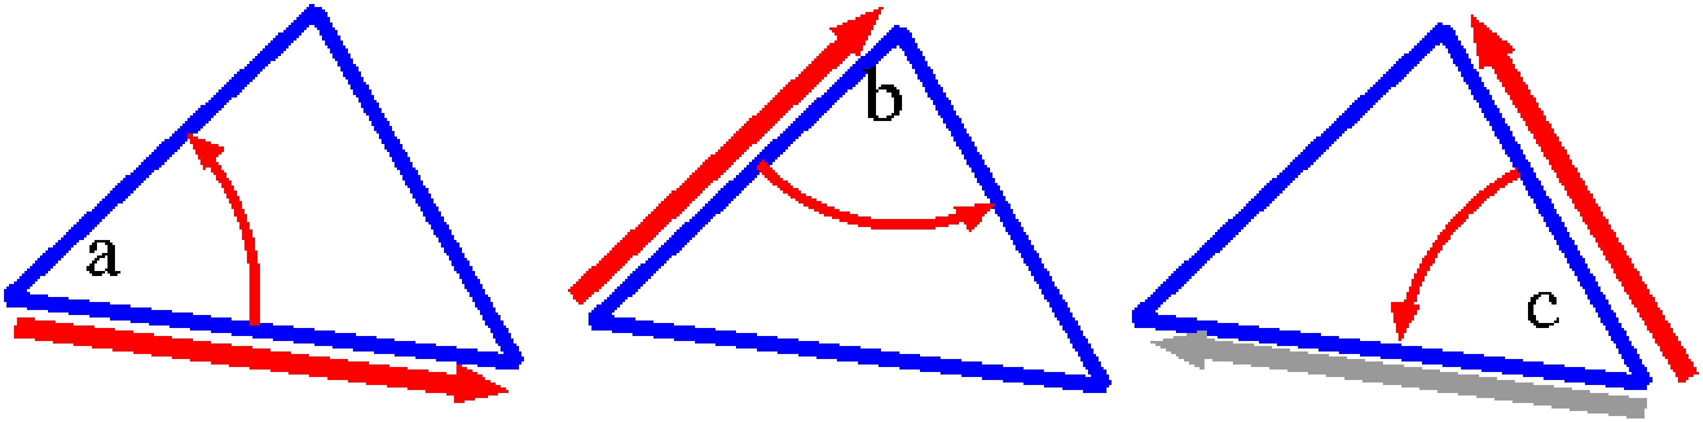

Examples of Euclidean reasoning 1a

This is just a tiny subset of examples: thousands more could be given.

Why do internal angles of a triangle add up to half a rotation?

A proof due to Mary Pardoe (when teaching maths about 40 years ago)

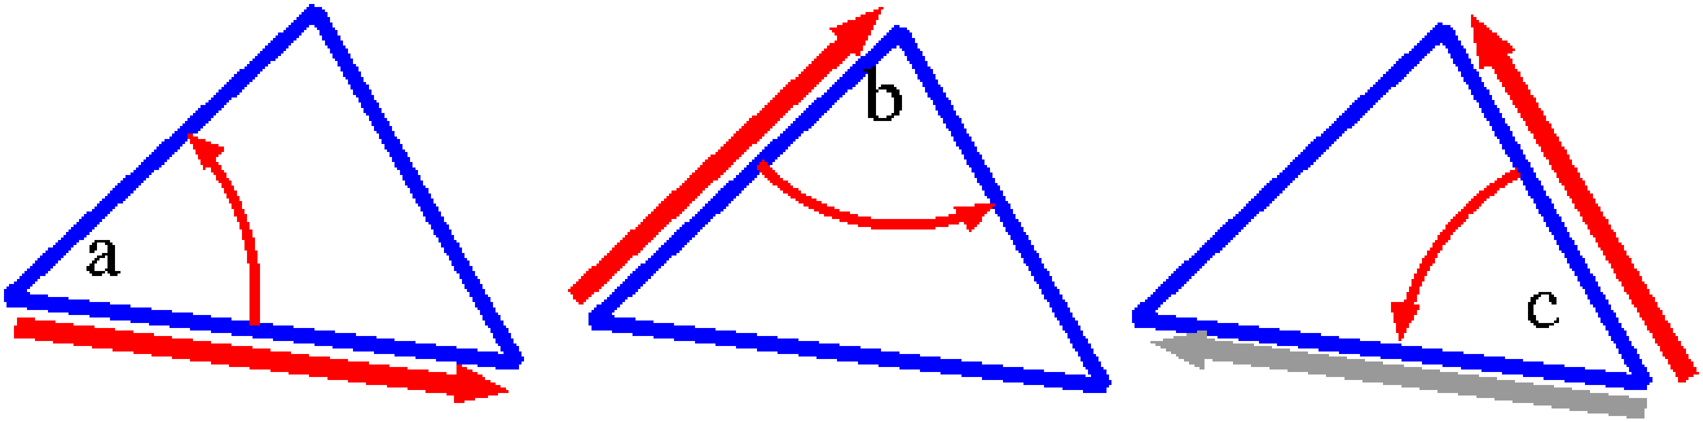

How do you know this proof does not depend on the size, colour, location, shape, orientation, ... of the triangle.

I.e. **it must work for any triangle?**

NB: this is not a form of probabilistic reasoning.

It's reasoning about an invariant property of a set of possibilities.

Mathematical necessity is NOT 100% probability!

Examples of Euclidean reasoning 1b

This is just a tiny subset of examples: thousands more could be given.

Why do internal angles of a triangle add up to half a rotation?

A proof due to Mary Pardoe (when teaching maths about 40 years ago)

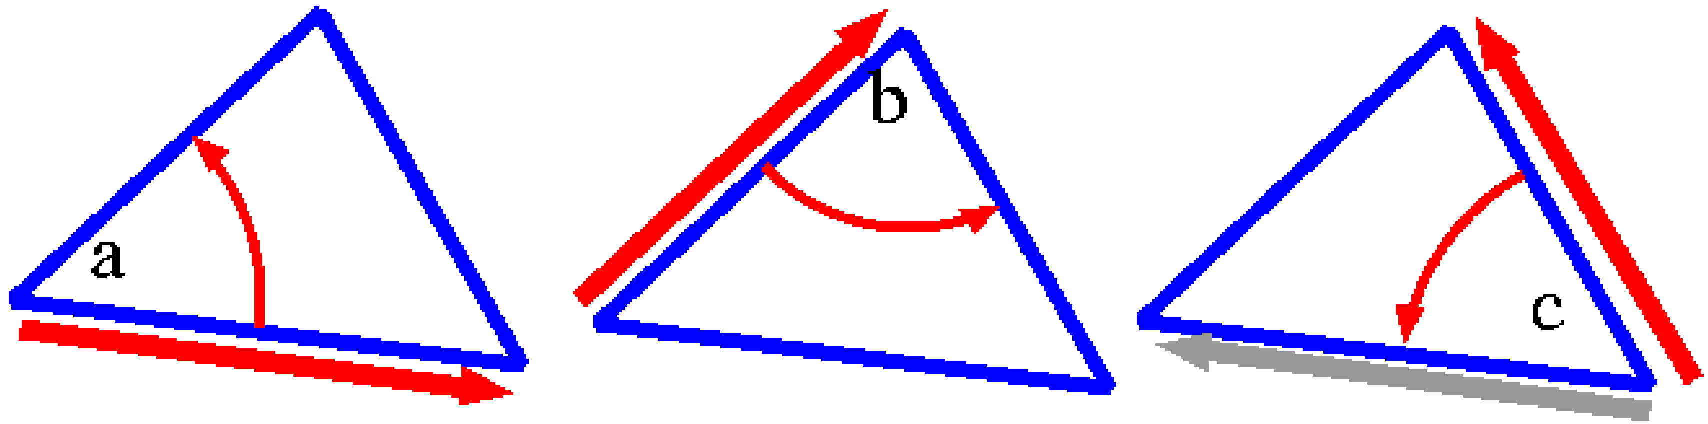

How do you know this proof does not depend on the size, colour, location, shape, orientation, ... of the triangle.

I.e. **it must work for any triangle?**

NB: this is not a form of probabilistic reasoning.

It's reasoning about an invariant property of a set of possibilities.

Mathematical necessity is NOT 100% probability!

Geometrical qualia have deep and rich implicit structure that mathematical information processing makes evident.

Examples of Euclidean reasoning 2

This is just a tiny subset of examples: thousands more could be given.

Reasoning about how area of a triangle changes if one side is fixed and the remaining vertex moves.

See this file for a discussion of some of the "hidden depths of triangle qualia" related to area:

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

Examples of Euclidean reasoning 3

This is just a tiny subset of examples: thousands more could be given.

“Chinese proof” of Pythagoras’ theorem (one of many):

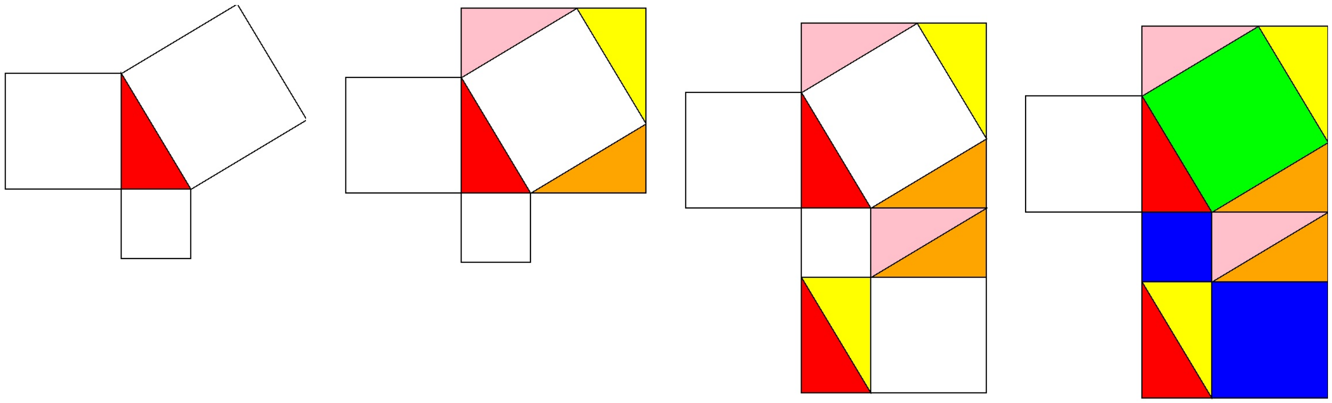

Start with a right-angled triangle and draw a square on each side (left figure).

The theorem states the large square (on hypotenuse) is the same area as the sum of the two smaller squares on the other two sides

Fig 2 shows three more coloured triangles all the same size and shape as the original red triangle, forming a big square, including the square on the hypotenuse.

Fig 3 shows four copies of the triangles moved to the picture below, and packed in a way that leaves two square spaces, the same sizes as the two small squares above.

Fig 4 shows that the big green square above must occupy the same total area as the two smaller blue squares below.

The online demo shows all the shapes changing in parallel, while relationships are preserved. Online demo based on an idea by Norman Foo

How can we know the proof doesn’t depend on the particular right-angled triangle chosen?

Main points 2a

To understand what evolved, why it evolved, how it evolved, what its current implications are, and how the lack of similar competences explains serious inadequacies in all current robots (that I know of),....

Main points 2b

To understand what evolved, why it evolved, how it evolved, what its current implications are, and how the lack of similar competences explains serious inadequacies in all current robots (that I know of),....

we need to think like **designers** starting with very simple organisms (or perhaps pre-biotic) structures and processes,...

Main points 2c

In order to understand what evolved, why it evolved, how it evolved, what its current implications are, and how the lack of similar competences explains serious inadequacies in all current robots (that I know of),....

we need to think like **designers** starting with very simple organisms (or perhaps pre-biotic) structures and processes,...

And try to understand far more of the changes in information processing that were required at various stages in evolution, in reproduction, in development, in learning, in social/cultural evolution/learning...

To extend all the research on evolutionary and developmental changes in: physical form, sensory-motor morphology, chemical bases of physiological processes, neural structures, physical environments, behaviours, competition, symbiosis, ...

(some of these overlap with changes in information processing).

Main points 2d

To understand what evolved, why it evolved, how it evolved, what its current implications are, and how the lack of similar competences explains serious inadequacies in all current robots (that I know of),....

we need to think like **designers** starting with very simple organisms (or perhaps pre-biotic) structures and processes,...

And try to understand far more of the changes in information processing that were required at various stages in evolution, in reproduction, in development, in learning, in social/cultural evolution/learning...

To extend all the research on evolutionary and developmental changes in: physical form, sensory-motor morphology, chemical bases of physiological processes, neural structures, physical environments, behaviours, competition, symbiosis, ...

(some of these overlap with changes in information processing).

This is "The Meta-Morphogenesis project", elaborated in

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/evolution-info-transitions.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/kinds-of-dynamical-system.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/pre-meta-config.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-requirements.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-sum.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/mathstuff.html
http://www.cs.bham.ac.uk/research/projects/cogaff/misc/beyond-modularity.html

What Alan Turing might have done.

The Meta-Morphogenesis project combines and extends the ideas in

- Turing's early (1936) work on Turing machines
  using discrete operations on discrete structures;
- His 1952 paper on Chemical Morphogenesis,
  using a mixture of discrete and continuous processes to explain striking global patterns.

See also

A.Sloman
Virtual Machinery and Evolution of Mind (Part 3)
Meta-Morphogenesis: Evolution of Information-Processing Machinery,

`http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d`
(Sloman, 2013)

Main points 3

Mathematics depends on the existence of domains of many kinds and sizes.

There are infinitely many of them because new ones can always be created out of old ones.

Some domains are instantiated in physical and chemical structures and processes (actual or idealised):

E.g.

- The systems of interacting point-masses of Newtonian mechanics
- The systems of interacting atoms and molecules in chemistry

  With an enormous variety of mixtures of discrete and continuous processes: e.g. catalytic operations altering chemical bonds, and molecules bending, twisting, coming together, moving apart....
- And many many more, some of them documented in
  `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html`

More abstract domains

Other domains are more abstract, e.g. the domain of one-to-one mappings and operations on them.

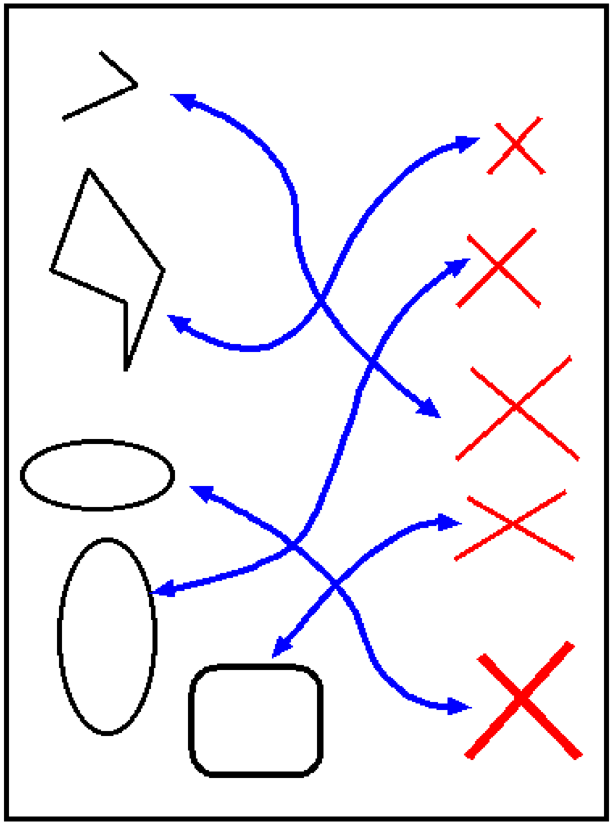

E.g. if you move around some of the crosses, will there always be one-one mappings with the shapes on the left?

Why doesn't the possibility of a one-one mapping between two collections depend on where the objects are?

Why are one-one mappings necessarily transitive?

Many domains involve information structures and processes

E.g. feedback control loops required to preserve some state.

E.g. feedback control loops required to achieve some state.

(E.g. processes of "hill climbing")

Life includes:

matter, energy and information, all interacting in very complex ways in

growth, metabolism, reproduction, then later on learning of various kinds, meta-cognition of various kinds, collaborative/competitive activities of many kinds.

(Ganti, 2003)

Evolution discovered many “theorems” about many domains

But it did so blindly:

The theorems were implicit in the solutions to problems of creating of structures, mechanisms, processes, etc. that had various useful (or harmful) properties.

The proofs were in the traces of evolution and development leading to those examples.

But often the proofs are wiped-out: e.g. there are no fossil records of information processing.

So we have an enormous amount of detective work to do to find out what sorts of intermediate stages there were.

And doing that could give us entirely new views about what now exists:

E.g. human mathematical competences that current AI systems and robots don’t have.

These are of central importance to many kinds of “offline intelligence” in humans and other animals.

The “blind watchmaker” metaphor focuses on only a small subset of the achievements of natural selection, compared with the “blind theorem-prover” metaphor.

We need to re-boot AI

A large proportion of the current AI/Robotics research, and a great deal of psychological and neuroscience research focuses only on “online intelligence”: walking, running, jumping, catching throwing, etc.... all of which can be achieved (in limited and inflexible ways) without offline intelligence – reflective intelligence, – mathematical intelligence.

It’s much too narrowly focused.

Later on the blind theorem prover produced

Later on the blind theorem prover produced

Sighted theorem-provers!

I.e. human reasoners with meta-cognitive competences able to inspect and reflect on aspects of their information processing and acquire insights into WHY things work as they do.

But there may be several steps in that process, some of them achieved in other species, e.g. corvids, elephants, squirrels, cetaceans, and apes, as well as humans.

We need to tease out the many intermediate steps to get a deep understanding of the later stages.

For example, I think a major step will be finding out how organisms represent spatial structures and relationships instead of just assuming it's all done as modern mathematicians and engineers do – using the arithmetic (Cartesian) translation of geometry.

In many contexts, the Cartesian translation adds spurious requirements for precision, often unattainable in practice, leading to tortuous uses of probabilistic methods, that would be unnecessary if instead of global metrics we used networks of partial orderings.

But that's a conjecture, and I don't yet know how to do it.

I am collecting evidence (on web pages).

Videos of children and animals

I have videos, especially of young children, suggesting that at least some of the time (I suspect much more of the time than anyone has realised) they are not doing the kinds of learning most people think is going on.

This depends on the power of “architecture-based” motivation, compared to “reward-based” motivation.

(I am not talking about intrinsic rewards.)

Instead of seeking rewards they are exploring ways of partitioning the world of structures and processes and types of information into domains.

Not because they are trying to achieve rewards, but because evolution selected mechanisms that respond to various situations by generating motives: to change, prevent, repeat, repeat with variations, preserve, alter, smash, etc.

Later they spontaneously inspect and reorganise what they have learnt: this is not a process of using evidence to modify probabilities.

The individuals cannot possibly know that such motives may help them (much later) produce more descendents.

Annette Karmiloff-Smith (Karmiloff-Smith, 1992) has extremely important things to say about this, which she describes in terms of acquiring expert behavioural competence then going through various processes of “representational redescription”.

For a partial discussion of how her ideas relate to what I’ve been saying see:

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

DEMOS

Child learning to crawl

Child with piano

Child with yogurt

Child with broom

Child with train

Child with train2

Several videos here:

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

Changes in information transfer across generations

Evolutionary transitions concerning modes of inheritance:

- changes in information transferred across generations,
- changes in how it is transferred
  (genetic information, in the broadest sense)
- precocial/altricial strategies, e.g.
- content transferred vs learning mechanisms transferred
- generic schemas to be instantiated transferred
- meta-schemas (to be instantiated to generic schemas) transferred,,,
- illustrated by Chappell-sloman-miall diagrams depicting "epigenetic layering" in the following slides, based on (Chappell & Sloman, 2007)

May be compared/contrasted with Waddington's "Epigenetic Landscape", depicted below.

Individual developmental trajectories I

Routes from genome to behaviour : the direct (original?) model.

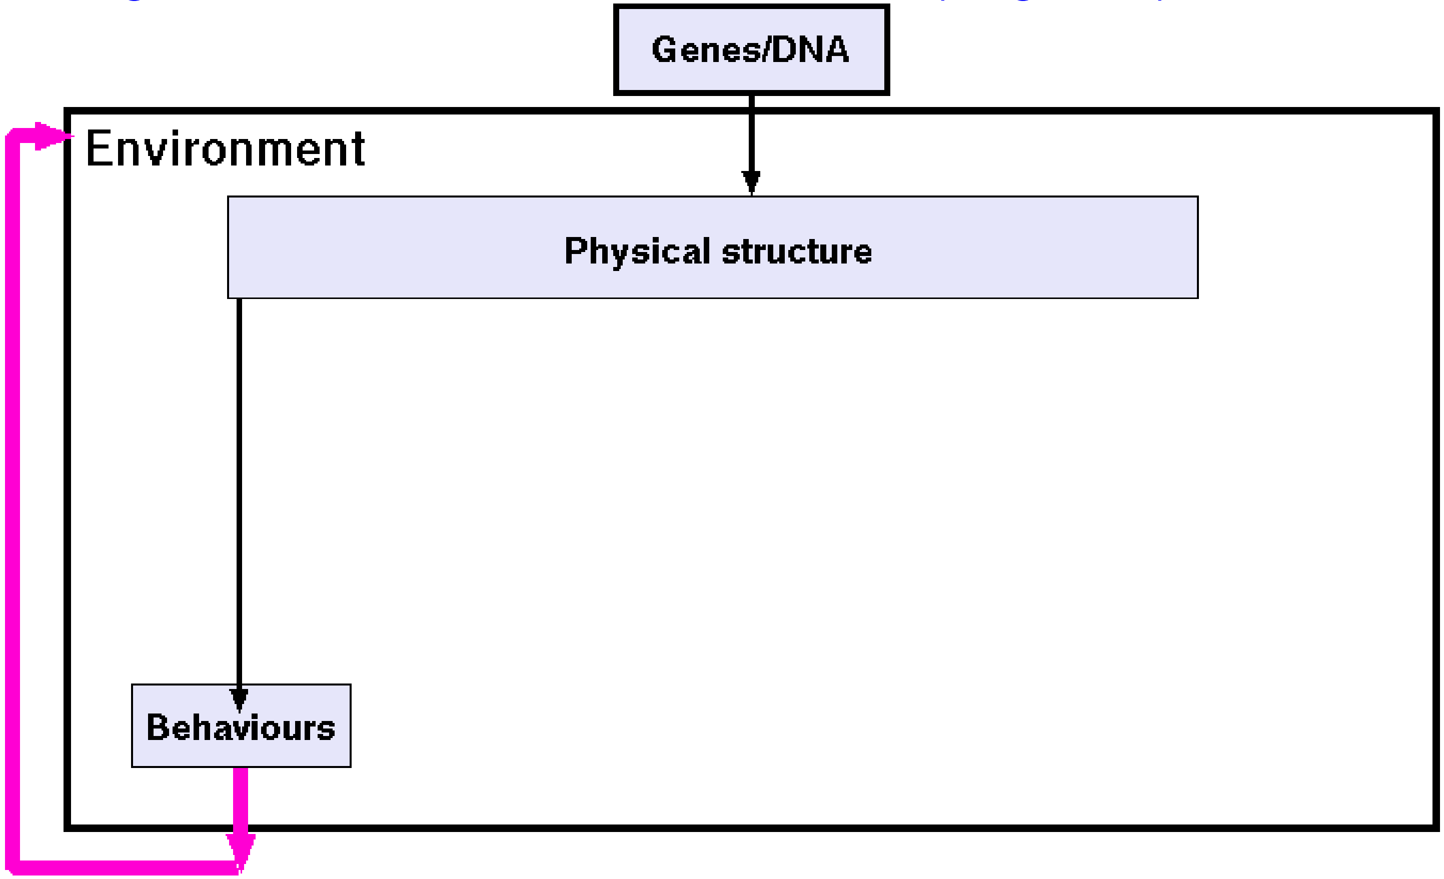

The vast majority of organisms (including micro-organisms) are like this.

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

Some of the more complex learning/development possibilities can be represented by alternative routes, as shown in the next few slides.

Individual developmental trajectories II

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

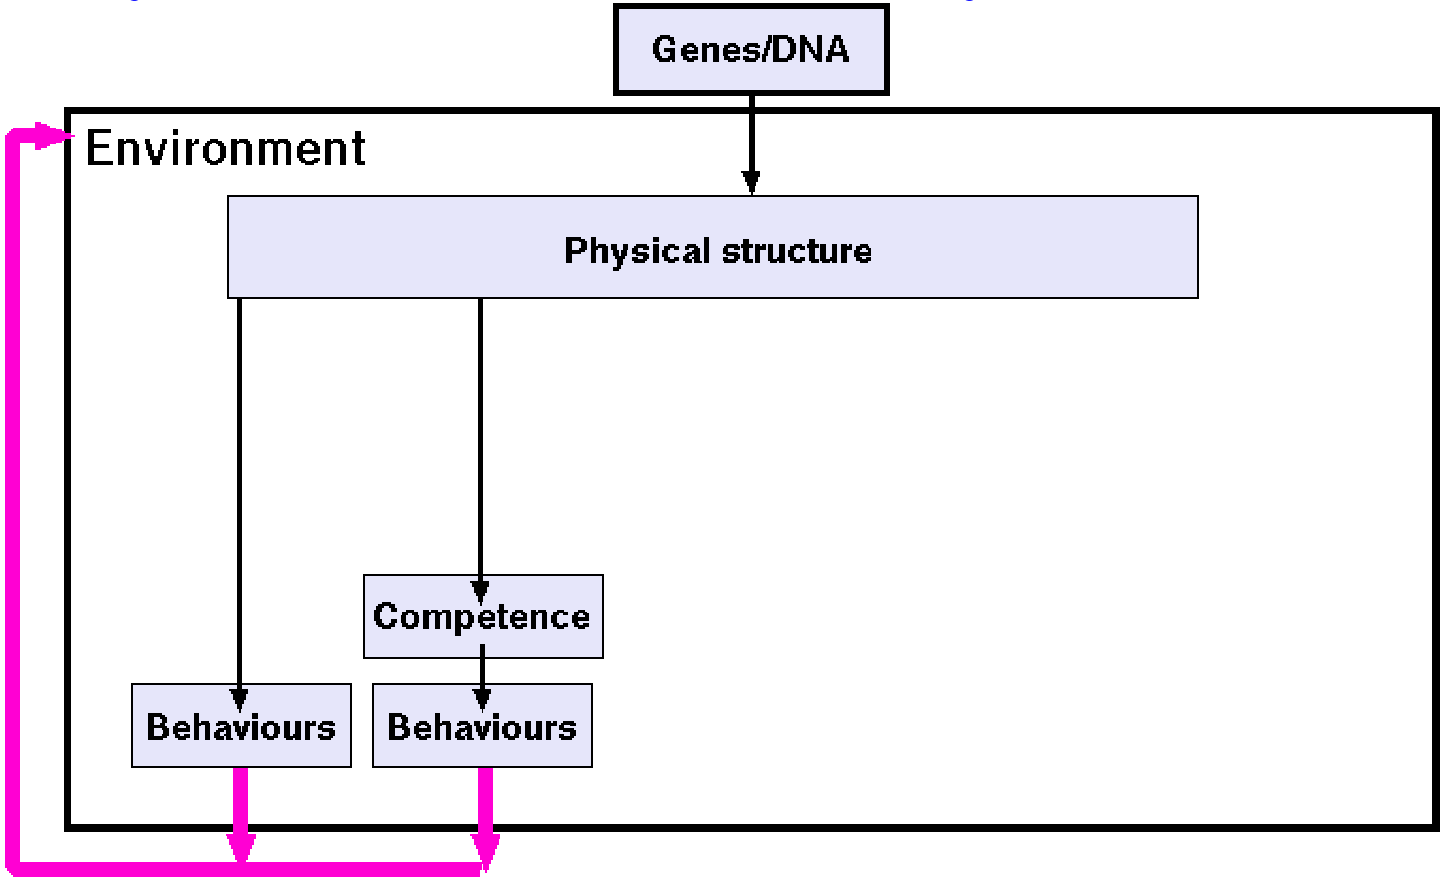

Some more complex organisms, instead of having only rigid (reflex) behaviours, also have competences that allow them to respond in fairly flexible ways to the environment: adapting behaviours to contexts:

reflexes trigger motives that produce behaviours instead of triggering behaviours directly.

Requires innate, or trainable, mechanisms for selecting means to achieve the goals.

Individual developmental trajectories III

Routes from genome to behaviour : stages added by learning.

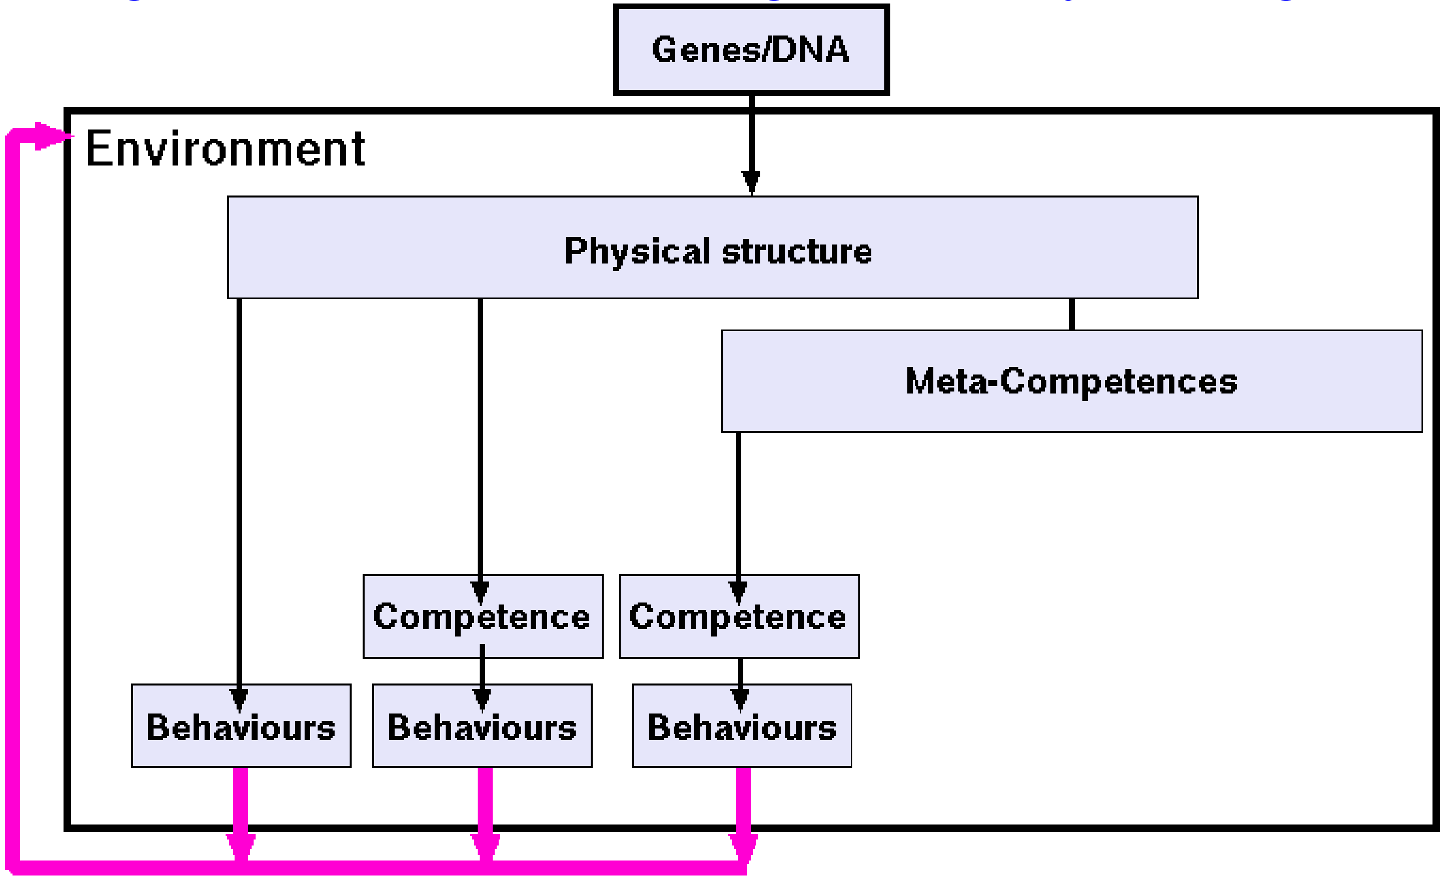

Genetically determined delayed meta-competences allow individuals to respond to the environment by producing new types of competence, based on learnt features of the environment, increasing flexibility and generality.

This requires extra learning mechanisms to allow appropriate competences to be generated and to allow those competences to select and execute behaviours.

Individual developmental trajectories IV

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

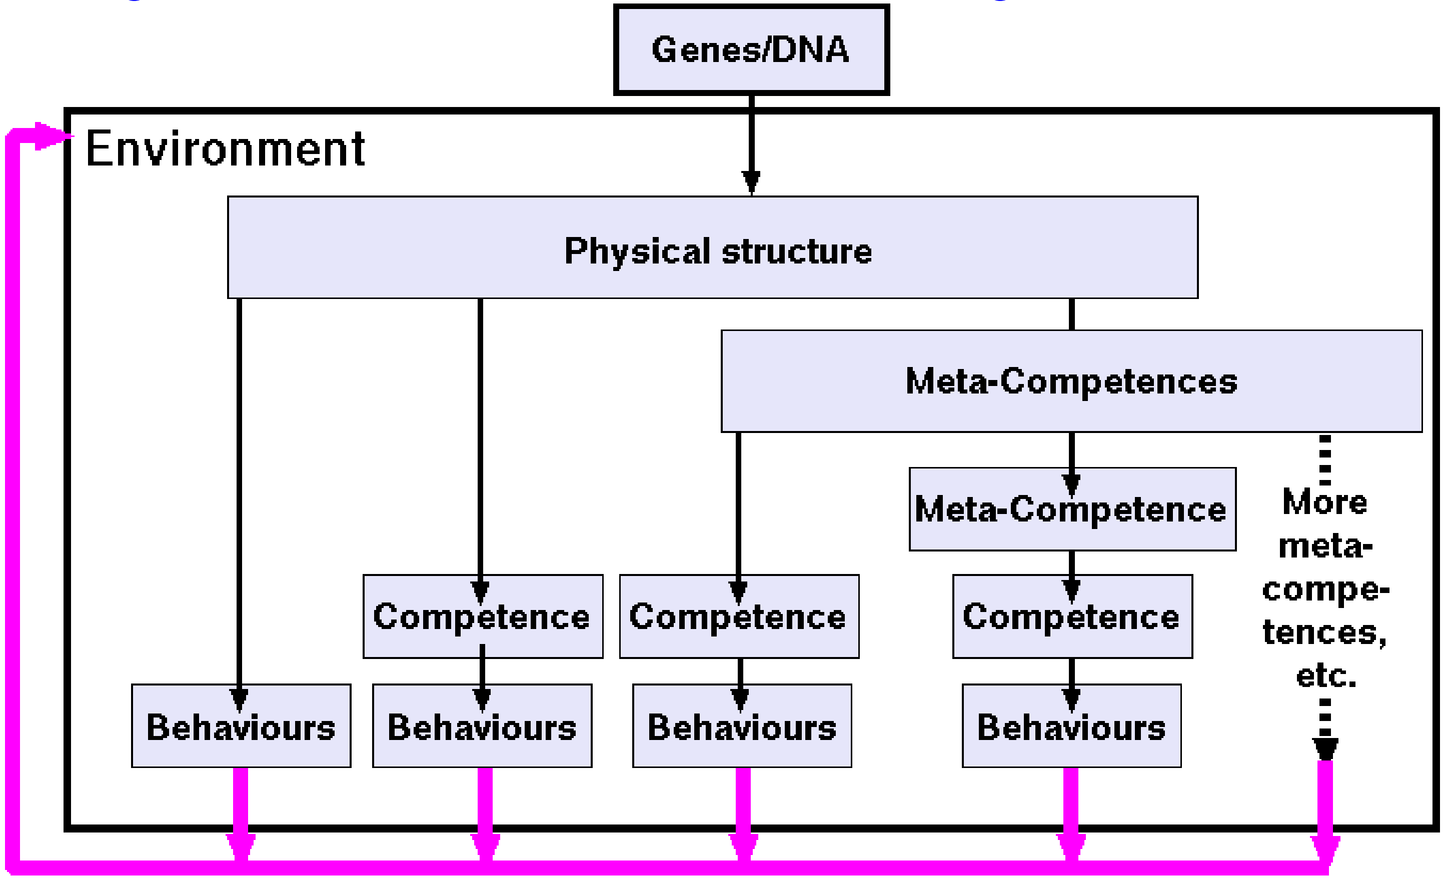

Some can also develop new meta-competences, on the basis of meta-meta competences.

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

E.g. learning to learn from reading, from complex experiments, etc.

(X) Compare Waddington's Epigenetic landscape

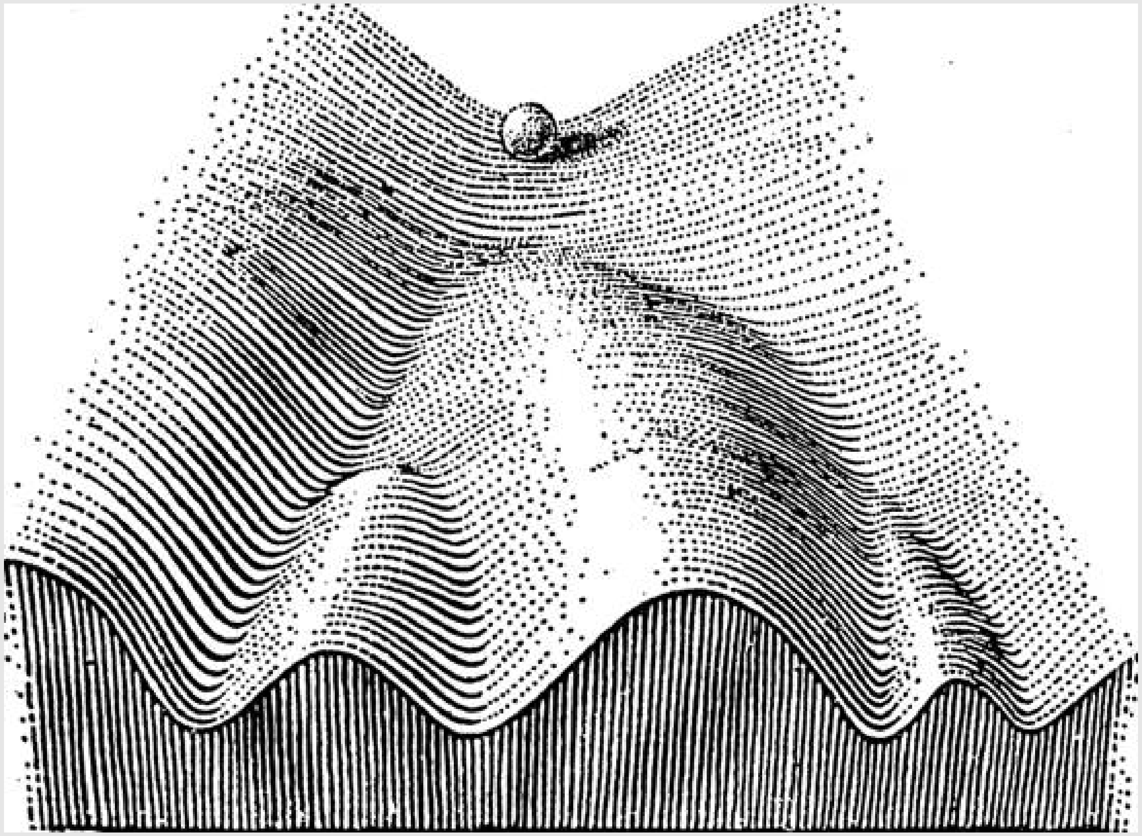

(Waddington, 1957)

Not nearly rich enough – I suspect he would have agreed.

Much more to be done

There are already many ideas about architectures, forms of representation, functions to be explained, generalisations of Gibson's notion of "affordance", varieties of motivation and "affect", new explanations of what qualia and how mechanisms for producing them might have evolved, followed later by mechanisms for discovering them (and misidentifying some of their properties!).

but also many gaps still to be filled.

I won't be applying for any funding: but I welcome interested collaborators.

There's lots more on the the web pages referred to, especially:

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

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

`http://www.cs.bham.ac.uk/research/projects/cogaff/misc/AREADME.html`

References

TO BE EXPANDED

References

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

Ganti, T. (2003). *The Principles of Life* (E. Szathmáry & J. Griesemer, Eds.). New York: OUP. (Translation of the 1971 Hungarian edition, with notes)

Karmiloff-Smith, A. (1992). *Beyond Modularity: A Developmental Perspective on Cognitive Science*. Cambridge, MA: MIT Press.

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. Available from `http://www.cs.bham.ac.uk/research/projects/cogaff/11.html#1106d`

Waddington, C. H. (1957). *The Strategy of the Genes*. MacMillan.

See also

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html

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

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/triangle-theorem.html

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html

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

http://tinyurl.com/CogMisc/evolution-info-transitions.html

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

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

http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html
Jane Austen's theory of information. (Long before Shannon.)