



Comments on Annette Karmiloff-Smith's (1992) book: Beyond Modularity: A Developmental Perspective on Cognitive Science

I now think this book should be essential reading for researchers in AI and Robotics as well as cognitive and developmental psychologists, educationalists, and philosophers interested in epistemology, philosophy of mathematics and nature-nurture issues, for reasons I try to explain below.

Table of Contents of the book

**Precis and commentaries in Behavioral and Brain Sciences Journal, Vol 17
No 4, 1994: DOI:10.1017/S0140525X00036621**

**Review by Michael Tomasello and Philippe Rochat (in Philosophical
Psychology 1993, partly discussed below).**

**K-S's work is also discussed by Margaret Boden in her Mind as Machine
e.g. pp 495--503 (Vol 1) and elsewhere.**

Annette Karmiloff-Smith's web page

Including recent publications developing some of the ideas in the book.

(DRAFT: Liable to change)

Aaron Sloman

School of Computer Science, University of Birmingham.
(Philosopher in a Computer Science department)

Installed: 13 Apr 2011

Last updated: 19 Apr 2011; 30 Apr 2011; 3 May 2011; 11 May 2011; 2 Jun 2011; 14 Jun 2011; 5 Jul 2011; 30 Jul 2011; 20 Oct 2011; 24 Mar 2012

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

It can also be accessed via this abbreviated URL: <http://tinyurl.com/bm-ks-as>

A slightly messy PDF version will be added from time to time (thanks to html2ps) here:

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

A partial index of discussion notes in this directory is in

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

Related presentations (e.g. on Piaget's ideas in his last two books, on necessity and possibility and on "toddler theorems") can be found in <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/>

CONTENTS

- [Background](#)
- [Toddler Theorems: added 22 Jul 2011](#)
- [Beyond Modularity -- an empirically-based theory of development](#)
- [Types of transition beyond "behavioural mastery"](#)
- [Draft more detailed list of of transition beyond "behavioural mastery"](#)
- [Conjecture: \(Added 30 Nov 2011\): transitions triggered by "density"](#)
- [Pre-configured and meta-configured \(precocial and altricial\) competences](#)
- [Where do the microdomains come from ?](#)
- [Other developments \(and Piaget\)](#)
- [Towards implementable theories](#)
- [No blooming buzzing confusion](#)
- [Relevance to Kant's Philosophy of Mathematics](#)
- [Humean and Kantian Causation \(Added: 30 Jul 2011\)](#)
- [Nature/Nurture issues](#)
- [The Discussion of Mathematics \(Child as Mathematician\) in Beyond Modularity](#)
- [The evolution and development of perceptual mechanisms](#)
- [BM Should be Compulsory Reading for Roboticists](#)
- [The Polyflap domain](#)
- [More things to read](#)

Background

I have known about this book for several years and even read [the 1994 precis](#) in Behavioral and Brain Sciences Journal, some years ago.

But for some reason, I thought I knew enough about the work to dispense with reading the book itself, though I occasionally referred to it in things I wrote, because I was aware of overlaps with ideas I was developing, e.g. related to the biological basis of mathematical competence in humans and creative problem solving in other animals, including work done in collaboration with [Jackie Chappell](#) on trade-offs between evolution and development (and the spectrum of precocial/altricial competences, which we now call pre-configured/meta-configured competences ([See below](#))).

Recently, following a tip from Karmiloff-Smith (**K-S**), I acquired and started to read Piaget's last two books, on Possibility and Necessity, because of their relevance to my research on development of cognition, as summarised here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/piaget-possibility-necessity.html>

I found the experiments reported by Piaget and his collaborators very interesting, but felt that the theoretical context in which they were presented was at the least very obscure, and possibly confused and inadequate because the theory was not based on an analysis of the information processing requirements for the behaviours observed. So I started offering what I regarded as a "rational reconstruction" of some of Piaget's ideas, in this presentation:

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

Piaget (and collaborators) on Possibility and Necessity -- And the relevance of/to AI/Robotics/mathematics (in biological evolution and development).

(My exposition is still very messy and incomplete. It will be improved from time to time.)

While searching for comments and reviews relating to those two books, and attempting to develop the rational reconstruction, I came across references to **Beyond Modularity, (BM)** and decided I should read it.

I now wish I had read **BM** when it first came out!

What follows is a very personal set of comments in which I try to explain why the book is so important, especially for roboticists trying to replicate animal competences, but also to philosophers of mathematics, as well as to developmental psychologists; and why it overlaps a lot with ideas I've been trying to articulate, arrived at from a very different direction.

In particular **BM** turned out to be very relevant to my attempts (begun around 1970) to understand the biological basis for the mechanisms that allow a human child to become a mathematician, which, if implemented in young human-like robots, could enable them to grow up to be mathematicians! (Some of the issues are discussed in 1978 in chapters 7 and 8 of The Computer Revolution in Philosophy (CRP78) -- based on earlier papers presented at AI conferences).

Toddler Theorems: added 22 Jul 2011

The key point that is not often noticed is this. Normally the only way in which humans now learn about mathematics appears to be by being taught. But at some earlier time in our history there were no mathematicians to do any teaching. Therefore what is now taught must originally have been learnt or discovered in some other way. The idea I have been trying to develop, e.g. in discussing "toddler theorems" is that there are deep evolutionary consequences of the fact that we and other animals evolved in 3-D environments that allow advantages to be gained from mathematical, or proto-mathematical abilities to reason about structures, processes, and interacting structures and processes.

See for example:

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

Online workshop paper (AISB 2010): <http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1001>

Example toddler theorem: a finger pointing downwards into a cup made of rigid material held in a vertical position can move from side to side beyond the limits of the cup. If the finger is moved down so that part of it is below the level of the top of the cup and it still points into the cup then only smaller sideways motion remain possible: larger sideways motions in that situation necessitate motion of the cup.

The key point is that after a certain kind of empirical learning combined with **reorganisation** of what has been learnt into a deductive system, the child can **work out** these specific facts from general principles, without having to experiment, and without being told or taught the "theorem".

It seems that during the first few years young children make many hundreds of such discoveries, some of them starting as empirical discoveries that later are transformed into results of spatial reasoning.

Other examples include facts about rigid levers (e.g. if pivoted along their length, then the ends must move in opposite directions), gear wheels (e.g. a pair of meshed gear wheels made of rigid material, able to rotate about fixed axles must rotate in opposite directions), inelastic strings (e.g. if a string is held at one end while the other end is fixed, and the held end is moved away from the fixed end, the string will become straighter, until it is straight after which the end cannot be pulled further from the fixed end. More complex toddler theorems about strings would allow for complications to do with other objects that can be involved, e.g. chair-legs, and also

complications involving knots.

After children start learning to count they can discover such facts as that counting a row of objects from left to right and counting the same objects from right to left must produce the same result. At first this is an empirical discovery, but later it is seen as a necessary truth, even though the child (and many adults) cannot explain why it is a necessary truth. (A theorem about one-to-one correspondences and permutations.)

There seem to be similar geometrical and topological reasoning capabilities in some other animals, e.g. nest-building birds, hunting mammals, elephants, and grey squirrels (based on observing them cope with "squirrel-proof" bird-feeders).

I am not saying that children (or other animals) that make such proto-mathematical discoveries are able to formulate them in English or any other human spoken language.

I think all of this is deeply connected with Immanuel Kant's claim that mathematical knowledge is both non-empirical (which does not mean it is innate) and necessary. (However, as Imre Lakatos showed, this does not imply that humans are infallible in their ability to make mathematical discoveries. Errors can occur, can be detected, and can be eliminated in various ways, depending on the type of error.)

There are many more examples in this wonderful little book by two authors partly inspired by Piaget. J. Sauvy, S. Sauvy, **The Child's Discovery of Space: From hopscotch to mazes -- an introduction to intuitive topology**, Penguin Education, 1974.

I have started compiling a web page on Toddler theorems here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/toddler-theorems.html>

I have recently been attempting a non-trivial conceptual reorganisation of my understanding of Euclidean geometry: attempting to see if there is a way of eliminating the parallel axiom, and instead extending the ontology to include time and motion, so as to allow an alternative proof of the triangle sum-theorem, discovered by Mary Pardoe, which for some people is easier to understand than the standard proof based on the parallel axiom. For details see this unfinished exploration:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/p-geometry.html>

All of this seems to come from evolutionary consequences of the fact that we and other animals evolved in a 3-D environment containing

- complex and varied kinds of stuff (material with different properties),
- structures of varying types and complexity (including many relationships between the parts of the structures and the parts of the parts, etc.)
- processes of varying types and complexity, in which relationships can change and interact
- discrete processes and structures whose elements can be linked physically or in thought or action. (Sloman 1978, Chapter 9).

In particular, an organism that can merely do "Humean learning" (even in its recent Bayesian incarnation), i.e. learning based on finding correlations (including conditional correlations) between experienced states and events will be severely constrained in its ability to deal with novel situations, e.g. involving new kinds of structures. Most species are bound by those constraints.

If, however, an organism can do "Kantian learning", it can reorganise what it has learnt empirically into a more structured generative, inference based system, which enables the learner to discover what is and is not possible, and what consequences of certain events will be, by *thinking* and *reasoning*. This requires information acquired empirically to be reorganised into something like a deductive system (an axiomatic system), though for non-human species or pre-verbal children the form of representation used cannot be human communicative language. However, I have also argued that some non-human species, and pre-verbal human children have perceptual, planning, and reasoning capabilities that make use of *internal* languages (forms of

representation) with the ability to support structural variability, compositional semantics, varying structural complexity, and inference. Jackie Chappell and I have called these Generalised Languages (GLs). See

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

Note added 28 Jul 2011

Sergio Pissanetzky (<http://scicontrols.com>) refers to the process of "Refactoring" in software engineering as a model for some evolutionary and developmental changes.

Structural Emergence in Partially Ordered Sets is the Key to Intelligence (2011)

<http://scicontrols.com/Publications/AGI2011.pdf>

When I eventually read **BM** I had already been thinking for many years (starting in my DPhil research presented in 1962) about this transformation of empirical knowledge into a more powerful proto-mathematical form, and my 2010 paper had conjectured that the "U-shaped" language learning observed in human children was a manifestation of a later evolutionary development of the older proto-mathematical biological mechanisms. This is closely related to K-S's claims about Representational Redescription and its cognitive benefits. She had reported the connection with language learning in 1992 (and possibly in earlier publications I have not read).

There are also connections with other research, which I shall not have time to spell out, including ideas in Minsky's two most recent books **The society of mind** (1987) and **The emotion machine** (2006), in both of which he emphasises multiple forms of representation and types of competence relating to the same domain, and also talks about A-brains doing things, B-brains monitoring and modulating what A-brains do, C-brains doing the same to B-brains, etc.

There's growing interest in Metacognition in AI which should also potentially overlap with this book, e.g.

Metareasoning: Thinking about thinking, Eds. Michael T. Cox and Anita Raja, MIT Press, 2011

However much of the AI work addresses problems of managing applied AI systems rather than problems of explaining or modelling animal minds.

Beyond Modularity -- an empirically-based theory of development (in humans and some other animals)

The ideas in **BM** are closely related to ideas I have been developing (partly with Jackie Chappell and with colleagues at Birmingham and elsewhere working on Cognitive Robotics), but Karmiloff-Smith adds a great deal of empirical detail and shows how development of a variety of different types of competence can be viewed in a uniform way as involving the transitions she describes as "Representational Redescription": transitions that occur in various domains after behavioural mastery has been achieved.

[

Unlike Piaget she does not claim that these transitions are related to age: the same sort of transition occurring in different domains, can occur at very different ages. Obviously there will also be differences in the mechanisms available to support and make use of those transitions at different ages. For example, an experienced mathematician developing competence in a field of mathematics may come to realise that there is a more abstract and economical way of summarising the field. Often the resulting process of abstraction produces a more general theory, applicable to other areas of mathematics.

Researchers who do not understand this process of abstraction and re-deployment, often confuse it with **metaphor**, a mistake because, unlike uses of metaphor, this process does not require details of the original domain to be preserved and used in future applications of the abstraction (like the 'source' of a metaphor).

]

These transitions are of different kinds, involving different mechanisms, and I suspect that an important subset of the mechanisms are also the biological basis of mathematical competences: they allow knowledge originally acquired empirically in some cases to be reorganised into a reasoning system in which some conclusions can be seen to have the status of theorems, e.g. theorems of topology, geometry, or arithmetic. Very young children seem to me to discover such things empirically at first, and then reconstrue them as non-empirical because they are derivable in what could be called a "framework theory". Only later do they become aware of what they have discovered, after using the discovery to work things out instead of depending only on empirical evidence. I have been trying to expand on that idea in various presentations here, including presentations on "toddler theorems"

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

I think this can be seen as an interpretation of Immanuel Kant's claims about some kinds of non-empirical discoveries being synthetic: they extend our knowledge. He does not claim that such knowledge is present at birth, since the mechanisms are "awakened" by experience. But what is then discovered is not empirical, just as mathematician exploring diagrams of various sorts can discover a theorem that is not empirical, but not innately known. K-S does not discuss this, but I think that the ideas are just below the surface in the book.

The theory in **BM** is a considerable advance both on Piaget's general ideas on development and also an advance on the presuppositions of many of the people working on machines (including robots) that perceive, act, develop, or learn -- because most of those researchers think only about how to get robots to achieve behavioural mastery (e.g. catching a ball, juggling, walking, running, following a human, picking things up, carrying things, going through doorways, avoiding obstacles, etc. etc., obeying instructions, answering simple questions, etc.).

Such behavioural mastery can be achieved without giving a machine the ability to think about what it has done, what it has not done, what it might have done, what the consequences would have been, what it could not have done, etc. Those additional competences require something like what K-S calls Representational Redescription, and we need to find ways to get robots to go through such processes if we wish to give them the kind of intelligence young humans, nest-building birds, hunting mammals, monkeys, and other primates seem to have. Developing such mechanisms will help us understand the processes that occur in children and other animals in a new, deep way. The book does not provide mechanisms, but K-S is clearly aware of the need to do so, and tentatively discusses some possible lines of development at the end of the book.

Her view on the importance of this is clear from this comment in her 1994 *Precis*:

"Decades of developmental research were wasted, in my view, because the focus was entirely on lowering the age at which children could perform a task successfully, without concern for how they processed the information."

Relevance to J.J. Gibson

Gibson's notion of perceiving and using an affordance is a special case of this general ability to think about what is and is not possible and why.

But Gibson's notion of affordance is too limited: it refers only to possibilities for and constraints on actions by the perceiver related to actual or possible goals or needs of the perceiver. I have discussed the need to generalise this (e.g. to cover proto-affordances, vicarious affordances, epistemic affordances, deliberative affordances...) in this book-chapter and in this presentation on vision research.

See also Piaget's books on Possibility and Necessity.

The idea of a "fully deliberative" system is also relevant

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

See also: A. Sloman, 'Actual Possibilities', in
Principles of Knowledge Representation and Reasoning: Proc. 5th Int. Conf. (KR '96),
Eds L.C. Aiello and S.C. Shapiro, Morgan Kaufmann, pp. 627--638, 1996,
<http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15>.

Types of transition beyond "behavioural mastery"

The main "high level" features of the transitions beyond behavioural mastery presented in **BM** seem to me to be right, though I think there are more different types of transition than the book presents, some of them specially relevant to developing mathematical or scientific understanding of a domain previously understood on the basis of empirical exploration and learning.

Note added 30 Nov 2011: I now see this as part of the Meta-Morphogenesis (MM) project:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>

The phenomena discussed in **BM**, and the evidence she presents are important independently of the precise details of the proposed taxonomy of types of transition and the conjectured explanations. (E.g. It can be argued that some of the transitions are architectural rather than representational.)

What follows is a draft list of kinds of transition leading up to and beyond behavioural mastery, transitions that should be distinguished in a more detailed analysis than K-S presents, though many of the ideas are already in the book.

I believe this list subsumes the list in **BM**, going beyond it in some ways. However, it may turn out that there's more in the book than I took in or remembered, so perhaps all the core ideas below are in it already. There may also be some important ideas in the book that I have forgotten to include here. (I have not yet done my second reading of the book. I wish it were available online, for easy searching.)

Draft more detailed list of transitions up to and beyond behavioural mastery

- Humans and some other animals acquire competences of many different kinds, which K-S describes as competences in "microdomains". (This is closely related to the idea of knowledge of "micro-worlds" in early AI.) Initially the knowledge about a domain is acquired empirically, e.g. by observation, exploration and experiment -- trying things out, often with minor variations, and finding out what happens.
- Later the knowledge about a domain (e.g. sucking things, grasping things, pulling things, pushing things, watching things move as they fall, roll, slide, bounce, etc) is transformed in various ways.
- Some new competences arise out of transitions from more primitive to more sophisticated forms of competence in the same domain, while others come from new combinations of old competences, as described in

- The most basic competence within a microdomain is partial or comprehensive "behavioural mastery", e.g. being able to walk steadily, being able to feed yourself with a spoon, or being able to communicate effectively about a range of wants, needs, games, toys and social activities.

Initially a competence may be unreliable and fragmentary, but learning, including much trial and error, along with accidental discoveries, guidance from older individuals, and observation of processes in the environment, gradually leads to increased behavioural competence.

Eventually there may be behavioural mastery (e.g. one-handed grasping and two-handed grasping) in a wide range of contexts, based on empirical learning and storage of particular behaviour fragments found to work for achieving certain goals in particular circumstances.

There may also be some limited ability to combine behaviour fragments learnt at different times, to produce novel behaviours, or to interpolate between learnt cases (something neural nets are good at).

- One of the most important kinds of transition (which includes various forms of Representational Redescription) is from empirically based behavioural mastery (mainly based on learnt associations, possibly encoded in Bayesian nets) to a more powerful, more generative, competence based on the discovery or creation (usually unconsciously) of new rules or principles, from which the old cases can be derived, but which go beyond the old cases, allowing new, better solutions to problems to be found and especially allowing solutions to be found to problems that have not been encountered before.

The best known example of this kind of transition (discussed in **BM**) is the transition, in language learning, from being able to produce and understand linguistic utterances composed of learnt fragments that can be combined in simple ways, to use of syntax-based comprehension and production, which provides far greater generative power.

The transition to the new more generative competence can go completely unnoticed by a young learner, even though it is evident to other individuals (e.g. parents, older siblings). When such transitions happen at a later age (e.g. developing mastery in chess, or programming, or mathematics, or musical composition) the learner is more likely to be aware of the transition (and indeed may explicitly aim for it) though often a teacher is needed to help the learner characterise what has and has not yet been learnt. Sometimes teachers who understand what is going on can provide both help and challenges to speed up and deepen the learning ("scaffolding") though for the most common forms of learning in infants and toddlers adults, and even developmental psychologists, are unaware of many of the kinds of changes in competence that occur. (Some are probably invisible at that stage, though they provide the basis for later, visible, changes in competence.)

This transition from early empirically based linguistic competence to a more generative syntax-based competence, can lead to new errors being made if the induced rules are incorrect, but children generally seem to correct such errors spontaneously. This initial success followed by partial failure, followed by new and greater success is often referred to as "U-shaped learning". Learners may make such transitions partly under the influence of genetically specified mechanisms, without having any awareness of the transition, though it can be detected by other older, more experienced, individuals.

Like K-S, I think this linguistic transition is a special case of something more general which evolved much earlier, and which occurs in children and other animals. Its non-linguistic manifestations are not so obvious since the products of learning are mostly not identifiable, recordable, structures, such as sentences.

Conjecture: (Added 30 Nov 2011): transitions triggered by "density"

One possible explanation for this sort of transition, which would be equally applicable to the linguistic and to non-linguistic examples, is that after the actions required to achieve various goals and prevent various unwanted states and events within a domain have achieved a certain "density" of coverage, some mechanism in the learner searches for recurring types of fragments of objects, fragments of processes, and recurring types of relationships between object fragments and process fragments. That can provide an ontology for collections of parts and relationships that cover the examples already acquired and can also be used to generate additional examples by applying the relationships to fragments that have not previously been encountered in those relationships, though they may have been encountered in other relationships.

This process induces something like a grammar for complex objects, situations, and processes.

Further data-mining processes can lead to discovery of generalisations about possible and impossible completions of fragments -- the discovery of "Domain laws". E.g. a motion of a large object towards a gap whose width is smaller than the object's cannot be completed by the object going through the gap. If the object can be part of a rotation which ends with one of its dimensions that's smaller than the gap being in the same plane as the gap, then the object can continue moving through the gap. **[Pictures needed]**

All of this may require creation of new forms of representation for the various structures, processes and constraints, and new mechanisms for manipulating those representations so as to reason about what is possible and what is impossible (i.e. its negation is necessary) in that domain.

These processes have to be repeated for many different domains, as illustrated, for example, in Piaget's books on Possibility and Necessity. This is referred to as "Acquisition of orthogonal re-combinable competences" here:

<http://www.cs.bham.ac.uk/research/projects/cosy/papers#dp0601>

There is much, much, more to be said about such processes and the mechanisms that can produce them.

See also: Aaron Sloman, If Learning Maths Requires a Teacher, Where did the First Teachers Come From?, in **Proc. International Symposium on Mathematical Practice and Cognition**, AISB 2010 Convention, Eds. Alison Pease, Markus Guhe and Alan Smaill, pp. 30--39, 2010, <http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#1001>

Compare the main idea presented in Kenneth Craik's 1943 book **The Nature of Explanation**, that some animals are able to construct and "run" models of things in the environment to make inferences, instead of attempting real actions to find out what will happen, or using only learnt associations.

It's an important idea, though so far the only work on implementing such competences uses a notion of running a model that is too close to simulating the actual process (as happens in game engines). That's different from what is required for reasoning about sets of possible actions with some common features.

AI planners, problem solvers, reasoners and theorem provers can be seen as special cases of this general idea, though most of the examples produced so far cannot be taken seriously as models of how human and animal minds work: they are too narrowly specialised for example.

There's a lot more to be said about the forms the new knowledge can take.

- A different sort of transition occurs when the learner acquires self knowledge about a competence. This can take different forms, including for example becoming aware of what the individual can and cannot do in various situations, or knowing why some action succeeds, or why doing it differently would produce different results.

A much more detailed breakdown is required.

- A further transition involves being able to communicate with others about the competence, e.g. explaining what to do and why, or providing advice, help, instruction.

A more primitive, less self-aware version of this does not involve explicit communication but it does involve changes in the type of confidence associated with actions, taking unhesitating decisions in certain contexts that would previously have require tentative trial steps, the rejection of misguided advice, the refusal to follow some bad suggestions, and so on.

- Yet another transition, which is closely related to my concern with the development of mathematics involves the ability to express knowledge about the domain in a systematic form of which one of the most advanced examples is the development of axioms of Euclidean geometry. Another example, which first happened centuries later is the axiomatisation of arithmetic. This is not a transition discussed in BM, but it seems to me to be the sort of thing that should be included in the list of transitions involving change of representation (along with change of information processing architecture).

This stage can have various forms, e.g. depending on whether the deductive system is fully formalised and in what way. The history and philosophy of mathematics includes several variants.

NOTE:

It is widely acknowledged that there are different kinds of numerical and quantitative competence, some involving discrete, others continuous measures of amount, some exact some approximate, some involving small sets others much larger sets, some involving notions of **amount** of stuff (e.g. water, sand), of area, of volume, or weight, of force etc.

However, the different types of mathematical sophistication required for these and whether or not perceptual or other mechanisms provide a basis for required total orderings, are issues that are not normally noticed. For example, the general notions of "amount of area", "amount of volume" and amount of matter in some object or portion of space cannot be comprehended without understanding something like integral calculus, which is not available to most young learners.

- There are also transitions in competences concerned with interactions with other agents. Animals, unlike inert physical objects such as piles of sand, pools of water, rocks, logs, are capable of self-initiated motion, including motion that achieves something for the individual, such as obtaining food, avoiding being eaten, finding a mate, protecting offspring, etc. A child (or non-human animal) may have some innate competences relating to other agents and some that are acquired empirically. For example, learning to cooperate with a parent while being fed, learning to prevent another individual getting a desired object, learning to catch an individual trying to escape, learning to avoid being caught, are all behavioural competences of types that are wide-spread among animals, though there are many detailed variations.

These competences require kinds of anticipation, collaboration, obstruction, and avoidance that are different from competences involved in interactions with inanimate objects. Let's call them interactive behavioural competences **IBCs**.

- There are kinds of change related specifically to **IBCs**, especially going beyond merely recognising and responding to visible behaviours to take account of internal states, including perceiving, intending, preferring, deciding, being afraid, learning, etc. This ability to represent and make use of information about internal cognitive and motivational process can add significant advantages, including greater accuracy of predictions, use of more sophisticated escape mechanisms, and also new useful forms of cooperation.

These can be called meta-semantic competences because they involve being able to refer to things that are able to refer. There are many debates about the mechanisms supporting such competences, including whether special forms of representation are required. I suspect that a major feature of this kind of development is architectural rather than representational, a possibility not considered in BM (unless I missed it).

There are well known philosophical problems associate with this kind of meta-semantic competence, including the need to be able to cope with opacity of reference.

- Sperber and others have analysed types of social competence, including communicative competences, that require meta-meta-semantic competences, or in some cases three, four or five levels of "meta-ness". There is much that is unknown about such multi-layered meta-semantic competence, and as far as I know little account has been taken of the detailed requirements for this sort of thing in work on robots that are intended to interact with humans.
- Another topic that needs to be addressed in more detail than BM does is the variety of ways in which ontologies referring to materials, objects, properties, relations, states, processes, and causal interactions in the environment can develop, initially through individual learning and development then later through socially cooperative development of scientific ontologies and theories. There has been some fragmentary work in AI attempting to model processes of ontology formation, e.g. by discovering invariants in mathematical relations. However this work is still in its infancy.
- Development of philosophical curiosity, including an interest in analysing complex concepts is another example, which takes many different special forms.

(... to be continued ...)

Some of the transitions seem to be driven in many learners by genetic mechanisms. Others, such as the explicit mathematical formalisation of prior knowledge, seems not to happen without the influence of cultural evolution, not discussed in BM. There are profound differences between

1. mere reorganisation of knowledge/competences from case-based (example-based) to a **generative** form that allows novel cases to be dealt with,
and
2. that reorganisation **coupled with** the ability to represent and reason about the change, and make use of information about the process of change e.g. to improve one's own competences or one's own learning processes, or to help another learner.

In all of this there appear to be deep connections with Immanuel Kant's philosophy.

Pre-configured and meta-configured (precocial and altricial) competences

It also seems to be clear that the mechanisms of change (learning, development, ontology extension, reorganisation of knowledge) are not all fully determined in the genome: individuals can learn to learn, or acquire new possibilities for development in ways that depend on personal environments and histories (e.g. going to university to study mathematics or philosophy).

These suggestions are consistent with the ideas in **BM** about the interplay between genome and environment during epigenesis. They are also closely related to the ideas in This paper written with Jackie Chappell. <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609> Natural and artificial meta-configured altricial information-processing systems (Invited paper in IJUC 2007)

about how some competences are preconfigured and others meta-configured, including this diagram (drawn with help from Chris Miall) showing how feedback from interaction with environment can significantly alter the late effects of influence of the genome:

-- XX

Where do the microdomains come from ?

Added 5 Jul 2011

The review by Tomasello and Rochat referenced above raises the question: "What is a domain?". Here's an extract:

Finally, we would just like to point out what Karmiloff-Smith already knows. A fundamental problem with the domain-specific approach is how one determines what is to count as a domain. Fodor has one way of doing it, Howard Gardner has a completely different way, some of the neo-nativists have other ways still, and Karmiloff-Smith has her own preferred way. Domain-specificity is a characteristic of human cognition in some sense, but it is not clear precisely how. It is our view, for example, that Karmiloff-Smith's attempt to establish written notation as a separate domain is not fully successful, and there are clearly in our view deeper connections between children's social cognition and language skills than she admits (and probably between objects and mathematics as well). And so, while we do not have answers to offer, we simply wish to caution that before the neo-nativist component of Karmiloff-Smith's model can be fully fleshed out, some theoretical work on what are the basic domains of cognitive development - and how we go about determining them - remains to be done.

I suspect there is no unitary notion of a domain that meets all the requirements, but we can make progress by characterising various special cases.

Suggestion: at least some of the microdomains are a result of processes of **carving up** the environment by means of classes of actions and classes of objects. Some of this process of domain-discovery seems to be driven by innate mechanisms. Others depend more on the environment, including the toys provided by adults. An example of such a domain, determined by a set of types of material and a set of

actions that can be performed on such materials is presented [here](#).

When children play there is often repetition with variations. I have noticed this in play with toys, with playground apparatus, with railings -- e.g. a child repeatedly swinging on a horizontal rail but with variations in the approach and the details of the movements, or putting a hand on a vertical pole and repeatedly going round it, again with minor variations. Different actions can be construed as being in different domains, but where the actions share substructures and share environmental objects and share temporal and spatial relationships, then it may be that the learning mechanisms use them to characterise a domain of competence that is at least initially separated from other domains. (Later, initially different domains can be merged through abstraction, through composition, or other processes.)

Some domains can be specified in a generative way, e.g. using something like a grammar, specifying types of element, ways of combining elements to form more complex elements, and ways of modifying already constructed elements. The "[polyflap domain](#)" designed to support research in robotics, is an example.

Further complexity can be added by specifying ways of increasing the variety of processes that can occur in the domain by combining processes into more complex processes. Computer programs are perhaps the most widely used way of generating very large and enormously varied classes of processes, but biological evolution (combined with individual development) may also have produced forms of representation and mechanisms able to perceive, generate, represent, reason about, very varied classes of physical processes.

The idea can be generalised to refer not just to classes of actions that the perceiver or some other intelligent agent can perform but to general classes of **processes** that can occur, whether produced by animal agency or something else -- gravity, the wind, impact of another object, earthquakes, etc.

One of the things that can happen after competences have been acquired in two or more domains is that the domains are **merged** to form larger domains. Sometimes this is just the sum of the previously existing knowledge, whereas in other cases the merging of domains adds something new. For instance a child who has learnt to play with water then learnt to play with sand has quite new opportunities when playing with water and sand simultaneously.

One of the features of the advance of science is the human ability to invent or discover, and then reason about domains of structures and processes that are not accessible to unaided perception and manipulation. An example is the development of theories of the atomic structure of matter and the processes that can occur as atoms are rearranged to form molecules of various sorts. A rather different sort of example is the development of Newtonian mechanics that required the use of an ontology including forces, locations, movements, accelerations, masses, presumed to have properties that can be expressed with mathematical precision going beyond anything that humans can perceive or measure. (This is a requirement for the application of differential and integral calculus in Newtonian mechanics.)

Other examples occur in other sciences. Human language is probably best conceived of as just one example among the many capabilities for coping with inanimate and animate entities in the environment exhibited by products of biological evolution, and increasingly, though **extremely** slowly being added to robots and other machines.

The review by Tomasello and Rochat also states:

Second, in the approach to human ontogeny proposed by Karmiloff-Smith, attention to the cultural dimension of development is minimal, as it is in Piaget. This is a serious omission, considering that a large part of human cognitive development consists in children learning from other human beings things that have been devised collaboratively by adults and modified to meet new exigencies over many generations of history. Thus the child does not need to invent the language that it learns to speak, the tools it learns to use, and one can only speculate what children's development of mathematical and notational skills would be if there were not some cultural conventions already existing....

This sort of comment ignores the fact that there were once no adult speakers. It follows from this that biological evolution must have produced the ability to **create** communicative languages, not just to **learn** them.

There is very strong evidence for this in the story of the deaf children in Nicaragua who, when brought together to be taught a sign language surprised everyone by collaboratively inventing their own, more sophisticated sign language.

See: http://en.wikipedia.org/wiki/Nicaraguan_Sign_Language

Of course, there are many examples where the ability of adult humans to teach and explain and the ability of young humans to learn from such teachers, makes it unnecessary for everything to be invented afresh by each generation. Instead there is a cumulative process. This applies to mathematical as well as linguistic competences, and many others. This ability to accelerate learning in such a way that each generation makes new progress by building on previous achievements seems to be unique to humans.

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

If Learning Maths Requires a Teacher, Where did the First Teachers Come From?

I have argued in the past that in other species, in our ancestors before human language evolved, and in young children who cannot yet talk, there must have been internal languages with some of the key properties of human communicative language (structural variability, unbounded structural complexity, context sensitive compositional semantics, and mechanisms for construction, interpretation and use of semantic contents). But such pre-verbal representational competences need not be fully specified by the genome: there can also be much individual development, influenced by the environment.

- <http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#713>
What About Their Internal Languages? (BBS Commentary)..
- <http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43>
The primacy of non-communicative language
- <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0702>

I believe this is consistent with the spirit of BM even though some of my conjectures go beyond what is in BM.

Other developments (and Piaget)

Piaget and others have studied kinds of development that involve motivation, values, preferences, policies, justifications, and affective states and processes, including enjoyment, wanting, disliking etc. These are topics for another occasion.

Unlike Piaget, K-S does not claim that each type of reorganisation happens only at a certain **global** stage of development -- the same type of development can happen at different ages, in response to learning in different domains (sensorimotor, linguistic, numerical, geometrical, pictorial, social, ...).

BM points out that what is discovered as a result of such learning is different in different domains, and typically depends partly on what the environment is like (from which the first level expertise comes empirically) and partly on innate and domain general developmental mechanisms that reorganise and reinterpret what has previously been learnt.

Her emphasis on the importance of the environment and the need for forms of development and learning that are tailored to specific sorts of environment seems to be correct, in contrast with all the researchers who seek a uniform, completely general, mode of learning that will work in any environment for a learner with no prior knowledge.

Compare Ulric Neisser:

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

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

John McCarthy made similar points in "The Well Designed Child"

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

However, she does claim that the processes that happen within a domain to produce such developments use general mechanisms, and I think that's not correct, because there is such a thing as learning to learn as discussed in the Chappell&Sloman 2007 paper.

But towards the end of **BM** she raises the same doubt herself (REF and QUOTE needed). Her self-criticism is one of the strengths of the book.

Towards implementable theories

For many in AI/Robotics and computer science the descriptions of mechanisms in **BM** will seem too imprecise or vague -- a problem discussed at the end of the book.

But it is better to get **the specification of what needs to be explained or modelled** right first, instead of taking some apparently obvious requirements (e.g. what **BM** refers to as mere "behavioural mastery") and attempting to satisfy them, as many AI/Robotics researchers do, ignoring important competences their machines lack, which go beyond behavioural mastery, including the ability to think about possibilities and constraints, and what might have happened but did not. (This is connected with the ability to perceive and reason about affordances.)

However, as the last two chapters of **BM** show, K-S is well aware of the need to go beyond her high-level descriptions and to use the "designer stance".

But the current concepts, formalisms and modelling tools in AI are not yet up to the job, for reasons that I have been trying to explain elsewhere. So it is important to go on identifying **requirements** for such tools, to help drive research in the right directions, and that is something her book does far better than most things I have read by psychologists about cognition. (I should produce a list of exceptions here!).

No blooming buzzing confusion

Many people agree with William James that a newborn infant has experiences that have little structure, only a "blooming buzzing confusion". I have argued here that that is a mistake. For something like human learning to occur there must, from birth or soon after, be fairly advanced perceptual systems geared to kinds of things that can occur in the environment -- pre-specified in the genome.

In **BM**, a similar argument is presented against Piaget, whose views about neonates seem to be very close to the idea of William James.

Relevance to Kant's Philosophy of Mathematics

It should be obvious that the ideas about representational redescription in **BM** are directly relevant to the task of trying to explicate Kant's philosophy of mathematics in terms of the biological ability to take some existing knowledge and then derive something new and significant from it *non-empirically*, which is the problem that first brought me into AI around 1971 (presented at [IJCAI 1971](#) in London).

The problem has turned out much harder than I thought 40 years ago, though I have recently been trying at least to specify some of the requirements for robot design that could produce a young mathematician (e.g. discovering simple theorems in geometry, topology, arithmetic, and seeing that they are not mere empirical generalisations (with or without probabilities attached) but necessary consequences of the assumptions. etc.).

I think most of what psychologists and neuroscientists have written about development of number competences is completely wrong, because based on an incorrect analysis of what numbers (of various kinds are). That's not surprising -- philosophy of mathematics is very hard.

Some of the requirements are presented in chapter 8 of [CRP 1978](#), including for example these prior competences:

- The ability to learn an arbitrary sequence of sounds or names
- The ability to recite that sequence
- The ability to perform other discrete actions, e.g. walking up steps, tapping objects on a tray, turning round, moving objects from one place to another.
- The ability to perform two discrete sequences of actions in parallel
- The ability to synchronise parallel actions
- The ability to use different stopping conditions for the same iterative process
- The ability to devise new procedures for solving new problems with this apparatus (e.g. answering "What comes before 'eight'").
- The ability to store partial results of exploration and use of the system, so that they don't need to be re-derived (though that raises problems about storage and access mechanisms).
- The ability to recursively apply these processes: e.g. to step through some or all of the sequence of names, in synchrony with stepping through another subset.
- The ability to replace the fixed set of learnt names with a generative mechanism.
- The ability to explore properties of the application of these competences and notice things (e.g. different orders of counting of a set produce the same result)
- The ability to **grasp** why that must be the case.
- The ability to explain to others why that must be the case.
- The ability to formalise in a systematic way the meta-knowledge acquired regarding a system of competences.

NOTE: [Heike Wiese](#) (Potsdam University) seems to have reached similar conclusions independently.

NOTE: most people do not get through all of these transitions in their development.

I don't know whether that's because of the different qualities of different teachers and educational environments, or because of genetic differences.

(**BM** discusses some influences of genetic abnormalities on cognitive development.)

See the presentations on "Toddler Theorems" [here](#).

Some of the overlap with Karmiloff-Smith's ideas is explained in my [abstract for an invited talk at AGI 2011](#)

An aspect of the overlap concerns nature-nurture issues and biological diversity. CRP78 presented (especially in [Chapter 8](#)) a challenge to over-simple theories about innate requirements for learning.

"...The old nature-nurture (heredity-environment) controversy is transformed by this sort of enquiry. The abilities required in order to make possible the kind of learning described here, for instance the ability to construct and manipulate stored symbols, build complex networks, use them to solve problems, analyse them to discover errors, modify them, etc., all these abilities are more complex and impressive than what is actually learnt about numbers! Where do these abilities come from? Could they conceivably be learnt during infancy without presupposing equally powerful symbolic abilities to make the learning possible? Maybe the much discussed ability to learn the grammar of natural languages (cf. Chomsky, 1965) is simply a special application of this more general ability? This question cannot be discussed usefully in our present ignorance about possible learning mechanisms...."

Humean and Kantian Causation

(Added: 30 Jul 2011)

One of many things David Hume is famous for is attempting to analyse the concept of cause and coming to the conclusion, based on his concept empiricism, that the only concept we can use is a notion of causation defined in terms of experienced correlation ("constant conjunction"), plus additional constraints especially temporal succession and spatio-temporal contiguity. He acknowledged that we often **feel** as if we are referring to something more, some sort of unobserved compulsion or necessitation linking the experienced states and events. (There is controversy as to whether he thought that feeling was an illusory subjective state, or something more. See http://en.wikipedia.org/wiki/David_Hume#Causation)

In response to Hume, Kant argued, (a) that concept empiricism must be false because we can't get all our concepts from experience because we can't have experiences without using concepts, and (b) that there are examples of causation that are not just instances of observed regularities, but involve a kind of necessity. E.g. changing the direction you go round a building changes the order in which you experience features of the building, and two rigid, meshed, gear-wheels must rotate in opposite directions.

I suspect he thought such examples of causation, which allow reasoning in novel situations by working out what **must** happen are very close to examples of mathematical necessity: e.g. certain changes of shape of a triangle cause the area to change, whereas others don't cause the area to change.

Jackie Chappell and I have argued that humans and other intelligent animals need both Kantian concepts of causation in order to work out **why** something has a certain effect when we have not experienced that effect previously (e.g. propagation of motion through a gear train), and also a Humean concept of causation in order to learn how to conjecture causes on the basis of observed correlations when they don't know why the correlations hold. Some organisms can use only Humean causation. Some may be born or hatched able to use only Humean causation, then later develop the ability to think about and make use of Kantian causation: humans appear to be an example, and possibly also Corvids and some primates among others. This transition in causal understanding may not happen at a specific developmental stage but can occur at different ages in relation to different domains of competence.

For more on these ideas see our WONAC workshop presentations: <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac>

How is this relevant to the ideas of **K-S** in **BM**?

I think one way of interpreting our claim about the transition from Humean to Kantian understanding of causation is as an example of what **K-S** calls "Representational Redescription". It requires replacing a collection of learnt generalisations about some domain of structures and processes with a **theory** of what's going on in the processes and why, which allows reasoning about causal relationships in novel examples of the domain. This can be seen as partly analogous to (though probably biologically older than) the kind of development in language learning from use of learnt verbal patterns to use of a generative syntax, mentioned above.

NB: the fact that a learner makes the transition to using Kantian causal reasoning in some domain does not imply that that transition is noticed or understood or explicitly describable by the learner. The development of meta-semantic concepts required to articulate what has been learnt may come much later. The earlier transitions may involve only the use of an internal, pre-verbal GL as described above.

Nature/Nurture issues

Before reading **BM**, I had been speculating about the fact that whereas biological evolution can produce highly sophisticated neonates (e.g. deer that run with the herd shortly after birth and ducklings and chicks that get themselves out of the egg, look for an adult to imprint on and also peck for food almost immediately) there are other species, including humans, that appear to have got very little from evolution, despite their adults appearing to be far more intelligent than the adults of species that start off more advanced.

Why doesn't evolution give the ones that need more sophisticated competences as adults a bigger boost at birth?

I learnt around 1998 (Thanks to Paul Marrow, then at BT) that the distinction I was discussing was labelled the Precocial/Altricial distinction by biologists.

Later, Jackie Chappell came to Birmingham and helped to clarify the issues. In particular our 2005 paper pointed out that the important altricial/precocial difference was a difference between competences, not whole species. E.g. humans have precocial competences at birth: sucking, breathing, etc.

We wrote two papers elaborating on some theoretical ideas about why the appearance of incompetence in newborn altricial animals was deceptive:

<http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>

The Altricial-Precocial Spectrum for Robots (Presented at IJCAI 2005)

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>

Natural and artificial meta-configured altricial information-processing systems (Invited paper in IJUC 2007)

Despite differences of emphasis there is a similar concern in **BM** with the diversity of biological learning and development across species, and also with an attempt to understand the evolution and functions of learning mechanisms that are very powerful, but not totally general, because they relate to features of the environment in which we evolved, albeit very general features.

An example of such a feature of our environment is that it contains 3-D spatial structures and relationships and processes in which the structures and relationships change, although the types of structures and processes can be very different in different places and at different times. (Compare a child growing up in a Cave-dwelling community with a child who is given lego blocks, meccano sets, and electronic construction kits to play with.)

Such evolved mechanisms, partly innately specified, can produce a wide variety of types of learning that occur at different stages of development. A feature of those mechanisms that is only implicit in **BM** is the use of extensions to the information processing architecture that develop at different stages in the individual, and which also evolved at different times. (They are all instances of the very generic (and still too vaguely specified) CogAff architecture schema)

BM also emphasises the fact that much learning that requires innate mechanisms can have features that are strongly dependent on what it is about the environment that is learnt (e.g. learning properties of physical structures and processes, learning to talk, learning about mental states and processes in oneself and others).

These can occur at different stages of development, and still be heavily influenced or constrained by the genome in addition to being influenced by the environment acted on and perceived.

She does not explicitly emphasise as we do (especially in the IJMC paper) that new learning competences can arise out of the interplay between the innate mechanisms (we call them meta-configured) and what has already been learnt. So a nine year old is capable of learning things a toddler cannot, and some of that may be dependent on having learnt ways of learning that are relevant to the environment: as shown by 21st Century nine year olds who have learnt how to learn to use a new type of software package, which none of their ancestors could ever do. Likewise a mid 20th Century physics graduate could learn new things about quantum mechanics that none of her ancestors had learnt.

This implies that cultural evolution, which is not much discussed in **BM** can feed information into the products of biological evolution that changes the learning and developmental competences produced by the genome. However this observation seems to me to wholly consistent with the general thesis of **BM**, and the discussion in **BM** of language learning, which has differences across cultures and across learning of spoken language and sign language, implicitly makes the point.

The Discussion of Mathematics (Child as Mathematician) in Beyond Modularity

Chapter 4 of the book "The Child as a Mathematician" requires critical discussion:

(a) because learning about numbers is a much more complex and diverse process than this book acknowledges (a tiny subset of the missing complexity is illustrated in Chapter 8 of CRP78, as explained above),

(b) because there is much more to mathematical development, including toddler mathematical development than learning about numbers, as indicated in presentations here.

Unlike **BM**, a central driver of my research has been a particular interest in trying to explain mathematical discovery, extending my attempt to defend Kant's philosophy of mathematics in my 1962 Oxford DPhil Thesis. Such knowledge (e.g. geometrical and topological knowledge) is clearly not empirical and also not restricted to innate knowledge, but rather, as Kant suggested, the discoveries are somehow *triggered* by experience even though they don't depend on what is experienced.

I.e. experiences can stimulate us to do things (certain kinds of reasoning) that do not depend on those experiences.

Chapter 7 of (CRP) on "Intuition and analogical reasoning", including reasoning with diagrams, and Chapter 8 "On Learning about Numbers" were specially relevant to the topics of **BM**.

Chapter 8 of CRP speculated about how a child who had already learnt to count could go on to learn new things about the counting process and about the individual numbers, including procedural discoveries (e.g. how to answer the question "What number comes before X?") and also factual discoveries that are then made rapidly accessible (e.g. discoveries about the properties of individual numbers that follow from their position in the sequence). I think those speculations are very closely related to, yet different in detail from, the theories in Beyond Modularity. (I would not use the label "Representational Redescription" because there's more going on than changes of representation, but I suspect that implicitly AK-S knows that.)

Because of the importance of external structures and notations in mathematical thinking, including the use of diagrams in geometrical proofs, I also emphasised the importance of the combination of innate mechanisms and environmental influences in building perceptual systems that can be used not only for perceiving what exists, but also for perceiving and reasoning about what is possible, and what is necessary -- the subject matter of Piaget's last two books. This is also a generalisation of J.J.Gibson's work on affordances, which I regard as a small but very important first step opening up largely uncharted explorations of the diversity of functions of perception, especially vision. (Chapter 9 of CRP was in part an attempt to show how different ontologies can collaborate in perceptual processes using extensions to the perceptual mechanism produced by learning. But the ideas were still rather simplistic and the importance of **perception of processes** was not noticed.)

(One of the features both of the discussion of use of diagrams in thinking -- which can either be diagrams in the mind or diagrams on paper, in sand, etc. -- and the discussion of architectures in chapter 6 emphasised that the mind/environment distinction needed to be blurred, since a blackboard or sheet of paper can be (temporarily) part of the mind of a mathematician -- as every mathematician knows. Google is now part of my long term memory when I am at my computer. This later become known as the "extended mind" thesis, discussed by Andy Clark, David Chalmers and others (independently of my work). [REF NEEDED].

I think some of **BM**'s discussions of children drawing can also be seen as supporting a kind of extended mind thesis, though only implicitly.

The evolution and development of perceptual mechanisms

Perception is not a topic BM discusses much -- though it is implicit in many of the experiments she describes, including drawing experiments. I think a computational investigation of the mechanisms required for both the perceptual and the drawing processes would significantly expand the ideas of BM, in part because there are special mechanisms required for vision that are different from most of the mechanisms she discusses, as I have tried to show in <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk88> and related papers.

In particular the perception of processes of varying kinds of complexity, both in seeing animal actions and in seeing mechanical interactions of inanimate objects, requires forms of representation that currently nobody understands, and cannot be found in AI/Robotic vision systems.

Addressing those problems (e.g. what's required to perceive the affordances in components of a meccano set while assembling a model crane as illustrated [here](#)) would add considerably to the content of this book, and also point to the need to go beyond her discussion of neural and other computational mechanisms in the last two chapters. (Remember, however, this book was written nearly 20 years ago.)

[To be continued]

There is a lot more to be said about the overlap, and our differences.

E.g. I would emphasise connections with Kant that are nowhere mentioned, and one of my major premises is that a test for a theory of this sort is that it can explain how humans make and systematise mathematical discoveries, and also how that ability evolved.

I think the ideas are deeply connected with AK-S' ideas about representational redescription, as illustrated in chapters 6, 7, 8, and 9 of my 1978 book: [The Computer Revolution in Philosophy](#).

It's interesting to see how work coming from such different directions overlaps, though it's also possible that I am hallucinating connections!

BM Should be Compulsory Reading for Roboticists

I now think **BM** should be compulsory reading for all researchers working in the areas of intelligent robotics, cognitive robotics, AI learning systems, and more generally computational cognitive science. Reading the book will give such researchers important new insights into what needs to be modelled and explained.

In particular, to the best of my knowledge robotics researchers focus only on the problem of producing behavioural competences: the lowest level of competence described by Karmiloff-Smith. They ignore, for example, an agent's ability to think about what it has not done, why it did not do it, what it might have done, what the consequences would have been, what might be going on out of sight, what possibilities for change exist in the environment, what the consequences of those changes would be, what new possibilities for change would occur if some change were made.

Robots that use (trained or explicitly programmed) simulators to predict consequences of processes or events, merely work out precisely what will happen in a very specific situation, but cannot think about why it happens and what difference it would have made if something different had been done, or if one of the objects had had a different shape or size or had been in a different location. It can only work out consequences for a totally specific situation, by running the simulation.

For example, a simulation program may be able to predict that if two co-planar gear wheels are meshed with fixed axles, then rotating one of them will cause the other to rotate in the opposite direction, but can only do this for the specific shapes and sizes of wheels in the simulation. Something more is required to abstract from the specific details and draw a general conclusion, or to reason about the consequences of the gear teeth having different lengths, or about the effects of the numbers of teeth in each wheel, or to explain how the situation changes if the teeth are not rigid. These additional competences are not mentioned in **BM**, but other cases are and I think there are many more cases illustrating the same points than the book presents. However, there are also important differences between the cases obscured by describing them all in terms of "Representational Redescription".

[To be continued]

I don't agree with everything in the book, in particular I believe that the emphasis only on **representational** change, without discussing required changes in information processing **architectures** during development, is a major gap (though she comes very close to discussing architectural changes, in discussing metacognition in Chapter 5). I shall gradually add constructively critical comments on the book to this page as I get time (and ideas). My own parallel ideas are expressed in a collection of presentations over several years, especially the presentations on vision, mathematical cognition, and evolution of language [here](#).

The Polyflap domain

The "Polyflap domain" developed a few years ago, and presented in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/polyflaps> (now used in the CogX project) was intended to provide an example domain in which adult humans can be learners like young children. I should now re-write that document relating it to Karmiloff-Smith's ideas in "Beyond Modularity".

The main point is that there is a lot to be learnt by playing with polyflaps, about how they look in various configurations how they can be arranged, how things change as new ones are added, and what differences are related to the polygonal shapes from which they are derived.

Some of that learning is empirical, but can be followed by a non-empirical reorganisation of what has been learnt into a deductive system.

The fact that the polyflap domain is unfamiliar to most people implies that it can be used to study child-like learning (RR) in adults. I think that will show up an extension that the RR theory needs: some of the kinds of reorganisation of knowledge discussed in BM can be triggered by when behavioural mastery is restricted and problems are discovered, instead of all RR following behavioural mastery. The book comes close to saying this, however.

More things to read

Some of the ideas about architectural growth and the tradeoffs between nature and nurture (or genome and environment) are presented in this paper

Jackie Chappell and Aaron Sloman, 2007,
Natural and artificial meta-configured altricial information-processing systems,
(invited paper) in
International Journal of Unconventional Computing, Vol 3, No 3, pp. 211--239,
<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>.

That was a sequel to this paper:

Aaron Sloman and Jackie Chappell, 2005,
The Altricial-Precocious Spectrum for Robots,
In *Proceedings IJCAI'05*, Edinburgh, pp. 1187--1192,
<http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>

Useful papers

<http://www.aaai.org/ojs/index.php/aimagazine/article/view/1041>
Oliver G. Selfridge, The Gardens of Learning A Vision for AI,
in *AI Magazine*, pp. 36--48, 14, 2, 1993,
Includes an example of a toy teaching model of learning in a simple
environment which illustrates a kind of representational redescription
that comes as a side effect of incremental learning. See the counting
example on page 44 and Figure 7. This inspired a Pop-11 program for teaching,
explained here
<http://www.cs.bham.ac.uk/research/projects/poplog/teach/finger>

<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.57.5839&rep=rep1&type=pdf>
Terry Dartnall, Redescription, Information and Access, in
Forms of representation: an interdisciplinary theme for cognitive science,

Karmiloff-Smith has a paper in the same collection:
Internal representations and external notations: a developmental perspective

TO BE CONTINUED

Topics for further discussion.

- The contrast between implicitly defined representations and explicitly defined representations.
 - How discrete domains are extended into continuous domains and how continuous domain are chunked into discrete domains.
 - How old domains are subsumed into new more general domains by abstraction, or composition, or both.
 - Different domain structures, e.g.
 - N dimensional spaces
 - trees, graphs, ... and their generalisations.
 - Logical/algebraic formulae or what they can represent
 - systems of relationships, including multi-strand relationships.
 - probabilistic/stochastic domains
 - domains that are rich in deducible theorems (toddler theorems)
 - domains with very little deducible content, only the possibility of empirical learning.
 - Domains with semantic content, requiring meta-semantic competences for their acquisition and use.
 - Domains concerned with self-monitoring, self-evaluation, self-control, ... meta-management.
 - Does BM allow for enough varieties of learning/discovery -- including learning by working things out, e.g. finding a proof of a geometric theorem or solving a differential equation, or discovering that positive integers include some prime numbers, and that there must be infinitely many prime numbers, since there can be no largest prime.
 - In BBS paper: "In other words, RR acts as a bridge for differently formatted pieces of information to be mapped on to one another." (p 734). Too vague.
 - In BBS: "If early number is based on magnitudes which do not map directly onto the counting system, then both must ultimately be in a similar representational format for a mapping between them to take place."
No: compare the relationships between integers and the real number continuum.
 - How can RR lead to **new** knowledge. Compare Kant and mathematical discovery.
-

Maintained by [Aaron Sloman](#)
[School of Computer Science](#)
[The University of Birmingham](#)