



Autistic Information Processing

Steps toward a generative theory of information-processing abnormalities.

(DRAFT: Liable to change)

Aaron Sloman

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

Installed: 23 Mar 2013

Last updated:

4 Dec 2014 (reformatted, slight improvements)

2 Jan 2014; 12 Apr 2014 [minor updates];

23 Mar 2013; 27 Mar 2013; 28 Mar 2013; 30 Mar 2013; 8 Apr 2013; 19 Apr 2013;

This paper is

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

A PDF version is available on request.

A partial index of discussion notes is in

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

TABLE OF CONTENTS

- [Abstract \(Added: 25 Apr 2013\)](#)
- [The need for a generative theory \(Updated 19 Apr 2013\)](#)
- [Summary so far](#)
- [The need for virtual machinery as well as physical/chemical machinery](#)
- [The need for meta-semantic competences](#)
- [Developmental trajectories in physical and virtual machinery](#)
- [Autism: A Different Sort of Information-Processing](#)
- [What does all this have to do with Nadia?](#)
- **Note added 30 Mar 2013:** [Low level visual processing is not image construction](#)
- [Developmental trajectories for "meta-configured" competences](#)
- [Waddington's epigenetic landscape idea](#)
- [The need for multi-level architectures](#)
- **Added 8 Apr 2013:** [Alarm mechanisms and varieties of emotion](#)
- [Requirements for explanations of cognitive abnormalities](#)
- [The need for a Meta-Morphogenesis project](#)
- (Added: 29 Mar 2013) [Notes on Motivation \(incomplete\)](#)
- **Note added: 2 Jan 2014**
["Are Prodigies Autistic?" by Scott Barry Kaufman](#)

- [Some references \(incomplete\)](#)
- [Related Papers and Presentations](#)
- [LIST OF FIGURES](#)

LIST OF FIGURES

- [Figure Word](#)
 - [Figure Layers](#)
 - [Figure Impossible](#)
 - [Figure EvoDevo](#)
 - [Figure Epigen](#)
 - [Figure CogArch](#)
 - [Figure CogArchAlarm](#)
-

Abstract (Added: 25 Apr 2013, revised 5 Dec 2014)

Trying to understand abnormalities of human minds leads to much scientific creativity. For example, if some some abnormality seems to be related to the lack of a competence, the absence of that competence may be thought to account for the abnormality -- a case being the suggestion that aspects of autism spectrum disorders are to be explained by absence of "theory of mind" capabilities, the ability to represent and reason about the beliefs and other mental states of other individuals.

However, there's not much point trying to explain a phenomenon in terms of non-development of something if you have no idea how that something normally develops and how it works (like trying to explain change blindness without any theory of what change detection is and how change detection works).

When the focus shifts to how some capabilities that normally develop work, and how the normal development can fail or be modified, and what the consequences of failure or abnormality are, we need a **theory** that explains how each phenomenon (e.g. development of an autistic disorder) relates to all the **other** things going on that can sometimes fail to develop normally, and whose interactions can produce cascaded effects.

That suggests the need to replace attempts to understand particular abnormalities (autism, Down syndrome, Williams syndrome, and others) with a **theory** about development and how development can vary, and why.

In that context the notions of "normality" and "abnormality" are scientifically relatively uninteresting, like the differences between common and uncommon chemical reactions.

Research on these issues can be distorted by focusing too much on particular syndromes and too much on the question "what goes wrong in this case", with too much emphasis on statistical correlations between **observed** results of development, instead of **hidden** (e.g. embryological, chemical, and neural) mechanisms producing those results.

I'll try to describe a methodological stance that focuses on types of **mechanisms** which may give us deeper insight and bring a wider variety of types of development under a common theoretical umbrella.

NOTE:

There are close connections with the ideas of neuro-developmental psychologist Annette Karmiloff-Smith in her 1992 book (Beyond Modularity) and her more recent work. A partial review of the book is [here](#). Some of my ideas have been deeply influenced by hearing her give lectures on developmental processes. More information about her work is [below](#).

The need for a generative theory (Updated 19 Apr 2013)

Capabilities or absence of capabilities in humans or other animals may require explanation. Different forms of explanation are possible, some shallower than others. Shallow explanations merely claim that the capabilities are predictable on the basis of previously observed correlations -- e.g. using characteristics of ancestors to predict or explain characteristics of descendants, or using previously observed correlations between infections or diet during pregnancy and characteristics of offspring.

In contrast, a deep explanation refers to mechanisms producing the characteristics. A type of mechanism is typically capable of producing a range of phenomena, and understanding how the mechanism works will make it possible to explain not just one set of characteristics but different possible sets of characteristics. In that sense an explanatory mechanism has "generative" power.

Normally human language capabilities are based on mechanisms with generative powers. That means that the mechanisms that explain how a person is able to say "I feel worried about my examination tomorrow" should also explain how the same person might have said "I feel confident about my examination tomorrow", "I feel worried about the election tomorrow", "I feel nervous about my presentation next week" and many more.

Similarly we should expect a mechanism-based explanation of how an individual came to have a particular type of abnormality also to be capable of explaining various kinds of abnormalities and also capable of explaining various kinds of "normal" development, since the same basic mechanisms of development are involved in all cases although details of their operation may be different either because of differences in the individual's genome, or differences in the context in which the genome is expressed (e.g. infection, abnormal diet, injury, behaviour of parents) or both.

A good theory of how a language works and how meanings are expressed should not be applicable only to particular specimens of that language. It should apply to all possible specimens of the language, including some that will never be uttered, such as possible variants of this sentence that extend it with illustrative phrases describing the appearances and behaviours of a variety of plants or animals. That requires a **generative** theory in the sense emphasised in Noam Chomsky's [early work on syntax](#), but not only a generative grammar that implicitly specifies an unbounded set of possible **sentences** but also a theory explaining how an unbounded set of possible **meanings** can be expressed in the language (a topic Chomsky did not address in that work -- I don't know about his later work).

More precisely, we need a generative theory of *intensional* semantics, not just *extensional* semantics (such as [Tarskian semantics](#)), since human meanings can refer to non-existent objects, processes, states of affairs, etc. and to objects whose existence is not required for the meaning to be expressed and understood. Without those intensional features of language we would not be able to ask questions about whether something exists, before knowing the answers, and we would not be able to have desires, intentions, and plans referring to possible future states of affairs and actions leading up to them, where possibility does not guarantee future actuality. [REFS to be added]

Note that a generative theory about meanings would have to include a theory about possible referents, and at least implicitly also about possible uses of language, e.g. answering questions, giving instructions, making inferences, explaining, and so on. This implies that a full theory about a language would have to include a theory about the world the language users live in, at some level of abstraction. This is important for understanding evolution of language capability, as well as the variety of languages in use, whether for communication or for internal use (perceiving, thinking, planning, etc.). This topic is important for a full theory, but will not be discussed further here.

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

for a discussion of functions and evolution of vision and language, including precursors of human spoken languages.

An even better theory should account for the possibilities of languages changing in various ways, including accretion of new forms of expression or new concepts, and also replacement of old ones.

Better still, the theory should apply to possible languages that have never been developed but might have been, just as the good theories of physics and chemistry apply to molecules that never have existed, and never will, but could possibly exist.

A still better theory would show how the mechanisms that allow individual humans to develop *linguistic* competences are a special case of more general mechanisms that account for a much wider range of competences, including perceptual competences, action competences, reasoning competences, the ability to have new kinds of motivation, the ability to acquire new kinds of knowledge, the ability to develop personalities, and many more.

Better still, the theory should be capable of explaining deviant or abnormal development, including both **outstanding** and **deficient** forms of cognitive development. The explanation could show which variations in the processes that are usually described as "normal" possibilities, can lead to known phenomena that are classed as "abnormal", and perhaps also explain how additional types of abnormality are possible even if they have never occurred, or have never been noticed.

Another requirement for a good theory is illustrated by ways in which other theories have improved. Humans once described and classified types of substance mainly in terms of how they look, feel, or taste to us, and how they can be seen to interact or made to interact with other things. (Perhaps that is the only way most animals can think about types of substance.) But we discovered later that a far more powerful and general theory is possible that classifies types of substance also in terms of how they are composed of chemical structures, that are composed of atoms, composed of sub-atomic particles, etc. So salt is composed of sodium chloride molecules and water is H₂O.

Diseases that were once classified in terms how they feel to sufferers and how the sufferers look and behave, are now more usefully also classified in terms of the causes of the diseases (micro-organisms, harmful chemicals, deficiencies in diet or environment, etc.) and the biological mechanisms interfered with by those causes.

This gives us a standard for a good theory of abnormal mental or development: it should also lead to a new, principled, *theory-based* classification of types of competence, forms of development, and types of abnormality, replacing classifications based only on the impressions of observers, the experiences of the individuals involved, or the performances observed in laboratory tests. I won't get that far in this document, but will suggest some small steps in that direction.

Summary so far

A good theory about the nature and origins of autism, or any other type of mental abnormality should not merely explain how that type can exist but should put that type and other known types of abnormality into the context of a theory that also explains how normal development is possible, and also explains the possibility of types of abnormality that are not yet known, but might be discovered, like molecules that are theoretically possible but have not yet been encountered.

Notice that this is not a requirement for prediction: a linguistic theory can explain how a particular lengthy sentence with a complicated meaning is possible in a language without providing any basis for predicting conditions under which such a sentence will be uttered. [Chapter 2](http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2) <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2> [Chapter 2](http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2) of [my 1978 book](http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2) explained the important scientific role of explanations of possibilities, and how meeting that requirement need not enable a theory to predict *when* those possibilities will be realised. That could be achieved by a later enrichment of the theory.

(According to Karl Popper, who was well aware that good scientific theories could have such a history, the earlier theory would be labelled "metaphysics" not "science", but I don't think such a classification is useful, given the huge importance for science of the earlier theories. See the aforementioned [Chapter 2](http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2).)

Likewise a theory that explains the possibility of a collection of human abnormalities may not be able to predict when those possibilities will be realised, perhaps because details of the physical mechanisms are not yet known or are too complex or too difficult to measure. A really good theory of this sort should be able to generalise to possible abnormalities in other animals, about which not enough detail is already known.

The remainder of this paper points out features of information processing architectures produced by biological evolution and features of the mechanisms of gene expression, and their complex dependence on the environment, that suggests a very complex, but systematically generated set of developmental trajectories, that should eventually explain the possibility of a wide range of cognitive abnormalities, including both exceptional talents and exceptional deficiencies. If it is a good theory, the possibilities explained will include hitherto unknown abnormalities (including exceptional talents and unfortunate deficiencies).

Compare: the atomic theory of chemical composition was a *generative* theory that explained the possibility of chemical compounds that had never been encountered.

The need for virtual machinery as well as physical/chemical machinery

Many biologists and neuroscientists would agree with the previous section, and regard it all as fairly obvious. But they may not agree with my next claim: that much of the developing machinery that needs to be understood in order to explain the phenomena in a new unified way, will turn out not to be neural or chemical machinery but virtual machinery, which is partly or wholly implemented in neural and chemical machinery (and sensory motor subsystems, along with features of the environment used by the machinery). Most of the important developments concerned with computers and information technology in the last half century have been concerned with increasingly complex virtual machines running on physical machines or on intermediate level virtual machines.

A brief introduction to this idea is available in a document on [Virtual Machine Functionalism](#), which includes pointers to additional background material. Some of the most important products of biological evolution have been developments of new kinds of virtual machinery, including the kinds

that develop in young animals as they interact with things in their environment, some of which are other members of their species, or members of other species. Examples of virtual machinery include perceptual mechanisms described below.

The need for meta-semantic competences

Some organisms are able, in their decision making, to take account not only of physical structures and processes in the environment, but also intentions, beliefs, plans, and capabilities to perform mental operations such as inferring, planning, choosing, preferring, and carrying out complex intentions. That requires not only the semantic competences mentioned previously as used to refer to objects, properties, relations, states and processes, but also meta-semantic competences, required for representing, thinking about, and reasoning about states and processes that use semantic competences.

Such meta-semantic competences may be used for self-reference (e.g. thinking about one's own past or future experiences or planning processes, and mistakes therein) or for other-reference (e.g. thinking about percepts, beliefs, intentions, plans, and mistakes made by others). The others may include prey, predators, or conspecifics, such as competitors, collaborators, mates or offspring who need help their thinking. Both self-directed meta-semantic competences and other-directed meta-semantic competences have been studied in AI as they arise in the problems of designing intelligent machines. Important early examples of self-directed meta-semantic competence in AI included programs that observed and improved on their own game-playing performance such as Arthur Samuel's [checkers player](#) developed in the 1950s and Gerry Sussman's 1973 [HACKER program](#) that examined its own thinking when its plans failed and learnt ways of avoiding similar mistakes at an earlier stage in planning, thereafter (using a meta-cognitive mechanism that he called a "critics gallery").

The ability to use meta-semantic competences in relation to others is often referred to as having "Theory of Mind" (TOM) competences. It is not always noticed by researchers that meta-semantic competences can be self-directed, or that meta-semantic competences require development of new kinds of information-processing machinery capable of referring simultaneously to things that refer and things they refer to, including cases where what is referred to does not exist -- e.g. imagined dangers producing real fear in a child. (For a general introduction to psychological research on TOM see [Apperly\(2010\)](#))

I have introduced the idea of a meta-semantic competence using words and concepts of our ordinary language for referring to human mental states and processes. However it is very likely that these competences are recent products of biological evolution building on a variety of intermediate competences that existed in our ancestors, many of which will be found in other species that interact cooperatively and competitively not only with physical objects but also with predators, prey and conspecifics. As in other cases of complex human abilities, e.g. visual perception, motor control, and various kinds of learning, it is very likely that current meta-semantic competences make use of and build on mechanisms supporting intermediate competences shared with other species. Some of those intermediate competences and mechanisms may play important roles during the boot-strapping of human minds from infancy, or before birth. As far as I know, nobody has developed a systematic taxonomy of types of semantic and meta-semantic competence and the intermediate cases required to account for the variety of known organisms that interact in more or less intelligent and informed ways with things in their environments. Research into such intermediate varieties of meta-semantic competence is an important subset of the Meta-Morphogenesis project summarised [here](#). See also [Fig EvoDevo below](#).

Without such a theory of how human minds make use of different sorts of virtual machinery with different sorts of semantic competence, including meta-semantic competences, we shall not be able to produce a generative theory of varieties of cognitive normality and abnormality. Some small steps in that direction are proposed below.

Developmental trajectories in physical and virtual machinery

Of course, the developments are physical also, including growth and change of sensory-motor organs, and many internal developments at all levels from the chemicals produced, through neural and other internal structures as well as physical powers such as strength, endurance and ability to digest and make use of new foods. If we don't understand all those patterns of development, and the ways they interact and the ways they can be perturbed, we cannot understand autism and other developmental disorders, including disorders with a common cause but divergent trajectories produced by environmental or genetic differences between individuals.

The work of Annette Karmiloff-Smith reported in her 1992 book, *Beyond Modularity: A Developmental Perspective on Cognitive Science*, and her more recent work on developmental neuropsychology is a pioneering attempt to produce such a generative theory.

I have some more detailed comments on her work, including her ideas about "Representational Redescription" in [this informal, incomplete, very personal, review](#).

Look also for her recorded lectures on development available online, e.g. M.I.N.D. Institute Lecture Series on Neurodevelopmental Disorders

- <http://www.youtube.com/watch?v=d2exgfkWWdc>
Disorders of Mind and Brain, uploaded on 5 Jun 2008
Makes the case for a cross syndrome, cross-domain, truly developmental way of examining human disorders using techniques like neuroimaging...
- <http://www.youtube.com/watch?v=Fb6BYeV9zIY>
Modules Genes and Evolution, Uploaded on 28 May 2008
Discusses insights gained from developmental disorders into how genes and evolution have shaped the brain and cognitive processing.

Researchers who agree with those claims may not all agree with my additional conjecture that the **information-processing products** of millions, or billions, of years of biological evolution are so complex that we have little chance of understanding how they work without attempting also to understand how many of their precursors worked. Very complex systems may depend in unobvious ways on previously developed mechanisms required for less complex systems, especially when complex virtual machinery is involved, whose components cannot be directly inspected physically (like the changing components of a complex piece of running software such as a chess program searching for a defence against a threat it has detected).

I am not endorsing the claim that "Ontogeny recapitulates phylogeny" explained and criticised [here](#).

That claim is usually taken to refer to transitions in bodily form and function in evolution and in individual development from a fertilised egg. I am referring to the possible continued use of many (possibly modified) very old forms of information-processing to underpin and provide some of the material for newer forms of information-processing.

The importance of understanding evolutionary as well as developmental transitions in information processing machinery, some of which is constituted by virtual machinery, is explained and illustrated in documents presenting [the Meta-Morphogenesis project](#) (work in progress -- building on earlier work on the [Cognition and Affect project](#)).

A good explanation of what goes on in cases of autism should be part of such a multi-layered, multi-strand developmental theory referring to multiple mechanisms that affect developmental trajectories. I'll try to illustrate this with a still very incomplete sketch that might be extended to incorporate various kinds of autism and other developmental abnormalities. However, far more work will need to be done. It may turn out that some of the work has already been done by others: I apologise for my ignorance and welcome corrections and pointers.

Autism: A Different Sort of Information-Processing

Many years ago, I learnt about Nadia, an autistic child who was able to make outstandingly good drawings from the age of three or four years onwards, but had very little linguistic competence and could not learn to read. Later she did acquire some ability to talk (I don't know details) and in the process apparently lost her outstanding drawing ability. She was described in:

Nadia -- A Case of Extraordinary Drawing Ability in an Autistic Child, by Lorna Selfe published by Academic Press, London: in 1977 Some of Nadia's drawings are in this online sample of a later book by the same author:

Nadia Revisited: A Longitudinal Study of an Autistic Savant, Psychology Press, 2011

<http://media.routledgeweb.com/pp/common/sample-chapters/9781848720381.pdf>

Pictures drawn by Nadia are now exhibited at the Gallery of the Bethlem Royal Hospital in Beckenham Kent:

<http://www.bethlemheritage.org.uk/gallery/pages/LD833-04.asp>

Selfe's book made a deep impression on me. I mentioned it briefly in [Chapter 8](#) (On Learning About Numbers) of [my 1978 book](#).

This document, however, is more concerned with implications of [Chapter 9](#), on multi-layered mechanisms required for visual perception, described below. Those mechanisms develop in normal children, and probably variants of them develop in many other animals, but apparently not in Nadia, with consequences described below. Building on suggestions made in Selfe's book, I'll outline a (partial) theory about Nadia's development below after presenting ideas about vision, and more general ideas about layered architectures, and then suggest possible connections with Nadia, and present a hypothesis that could be tested by a research programme investigating normal and abnormal developmental trajectories, of many kinds.

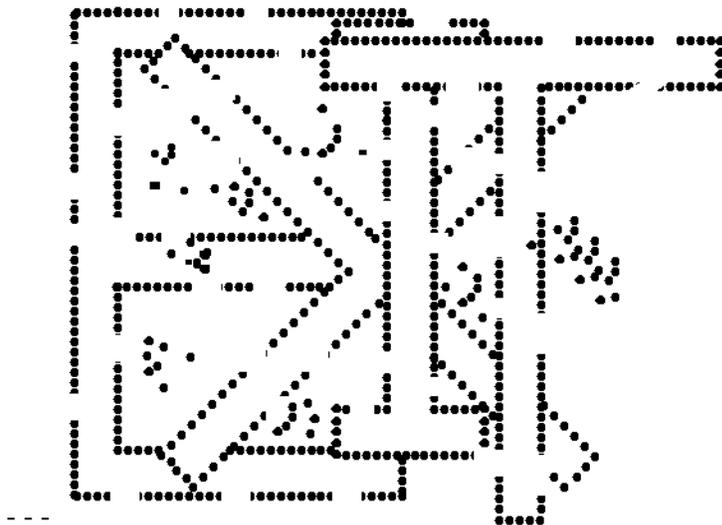
Thinking about Nadia in the 1980s, I noticed possible connections with the (partial) theory about architectural requirements for human visual perception mechanisms, developed with colleagues at Sussex University (David Owen, Frank O'Gorman and Geoffrey Hinton) and implemented in a "toy" demonstration program called "POPEYE" (because it was implemented in the Edinburgh AI programming language POP-2). The theory, and aspects of the program, were described in Chapter 9 of the 1978 book, available [online](#).

I shall present only a very compressed summary here.

The Popeye program could analyse and interpret pictures made of dots, where the dots formed straight lines, parallel pairs of straight lines formed "bars", and bars could represent flat plates some of them connected to others at "bar junctions", or partly concealing others, and connected flat plates represented capital letters, which together formed a word, in a previously known list of words.

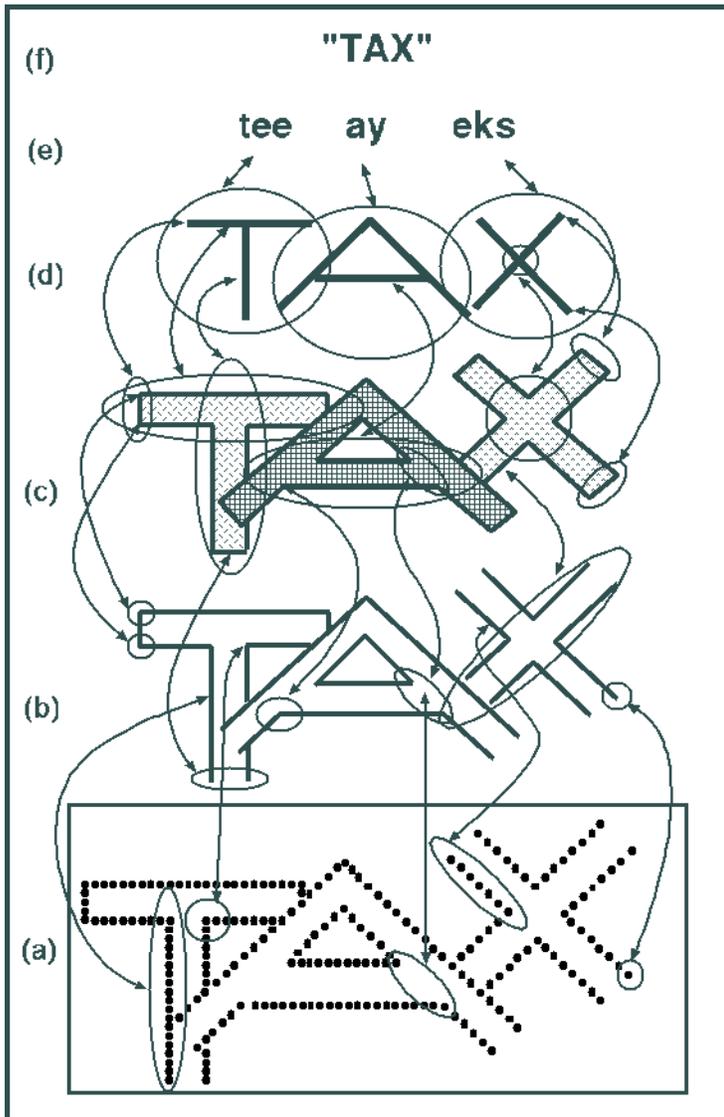
The task of the program could be varied from easy, with the letters clearly separated, to very difficult, with the letters jumbled together in various ways and with different amounts of positive and negative noise added. A fairly difficult case is shown below in **Fig Word**.

Figure Word



The next figure **Fig Layers** indicates the different layers of abstraction used by Popeye, with information flowing in parallel, bottom up, top down and sideways within layers, in such a way that Popeye was often able, like humans, to identify a word before all the lower level sub-systems had completed their processing.

Figure Layers



[Arrows crudely and incompletely represent both information flow and control flow.]

That mechanism was also able to identify words in noisy and jumbled pictures such as **Fig Word** above, which would have defeated most "pipeline" systems requiring analysis by lower layers to be completed before more abstract layers began their processing. However, Popeye did not understand English, did not handle sentences, and could not make any use of non-pictorial context. It also could not deal with curved lines: it was intended only as a proof-of-concept prototype, not an accurate model of human visual processing. (Related ideas have been developed by other researchers in vision.)

The Popeye project directly contradicted claims that were being made at the time by some well known vision researchers (which may have explained why our request for funding to continue the work was turned down).

However, similar ideas about mechanisms required for understanding human spoken language were being developed by researchers in the DARPA speech understanding project in the USA, e.g. the HEARSAY project, described in "The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty" by L. D. Erman, F. Hayes-Roth, V. R. Lesser and D. Raj Reddy, available here:

ftp://shelob.cs.umass.edu/pub/Erman_Hearsay80.pdf

One of the important features of Popeye was that it was not composed of a single algorithm that takes in an image and goes through a sequence of programmed steps to reach a conclusion about the word depicted. Instead it had a varied collection of sub-systems working in parallel on whatever information was available to them. The lowest level mechanisms could get information only from the image, or from the results of similar low level processes examining other parts of the image, or possibly via hints from higher levels suggesting that a line fragment might be present where none had been detected, causing a lower level line-detector to lower its threshold for line-evidence. The other levels could not directly access the image information but could use information from several other subsystems operating at different levels or on different parts of the interpretation at the same level.

A consequence of this architecture was that various mechanisms could be turned off, or made faulty without the whole system collapsing. We did not systematically investigate this or demonstrate it, because at the time our interest was not concerned with effects of malfunctions. In the present context the relevance is that if the processes constructing minds and brains also create systems with relatively independent sub-systems, then many different types of abnormality might be produced by damaging or removing different subsystems. Likewise super-normal functionality could be produced by adding a new mechanism, or collection of mechanisms, that add to the competence of the system -- for example asking someone else to suggest an interpretation for a part of the image, or overhearing someone else talking about the contents of the image and using the information in one of the overheard utterances to resolve an ambiguity in visual information found so far.

These comments about varieties of functionality help to illustrate some of the variation possible in results of multi-stage developmental processes discussed below, that can take different directions depending on the environment or on the influences or non-influences of other mechanisms produced by the genome, or the influences of previous kinds of learning. The set of possible outcomes of such a system could be so varied that even the word "spectrum" sometimes used to describe varieties of Autism or other abnormality, understates the richness of the possibilities for development inherent in each individual.

[NOTE

These ideas about vision, partly based on inspiration from Max Clowes after he came to Sussex university in 1969 (see [this tribute](#) for more information about him) are still under development, though mostly ignored by researchers in human, animal and machine (AI/Robotics) vision, and brain researchers. See [this incomplete discussion](#). **ENDNOTE]**

What does all this have to do with Nadia?

Selfe conjectured that brain damage had prevented the development of normal higher-level processing in Nadia of a kind required for language, and this somehow facilitated compensatory development of other capabilities. So, if something had prevented normal development of the more abstract layers of a visual information-processing architecture such as Popeye's, including perhaps the ability to develop abstract classifications of types of object, types of property, types of relationship, types of process, and types of action -- except for a few special cases, including those related to her fascination with making pictures -- then perhaps the neural resources that would have been recruited for more normal development were instead used to extend and enrich the mechanisms for analysing and recording low level image details. See the speculation in section 11 of [my 1989 paper on vision](#).

So, perhaps her visual perceptual mechanisms (and others) could not provide the semantic contents required for understanding a human language. The use of written and spoken language also requires abilities to "chunk" perceptual information into types that correspond to linguistic components, such as syllables, words, phrases and sentences, and in the case of written language letters and words. Use of communicative language also requires meta-semantic competences, including the ability to treat others as having thoughts, questions, intentions, etc.

It has been conjectured by many researchers that a feature common to Autism is underdeveloped "Theory of mind". From our point of view that is just a *subset* of what may be under-developed. There could be different patterns of partial development or abnormal development in different individuals labelled Autistic.

The non-development, or very partial development of some of the more abstract processing layers, would have left some available processing resources (e.g. neurons) unused that would normally be shared between layers, or dedicated to higher level layers. In Nadia, it seemed that far more processing resources than normal were focused entirely on extracting detailed "low level" information about the structures and relationships of local features of the visual input, such as 2-D relationships of image features. That is exactly the kind of information required for producing drawings that accurately represent the appearances of things. So the availability of a very powerful, highly developed low level visual processing system using a lot of brain power normally allocated to other tasks, might explain Nadia's ability to produce drawings that most people found difficult to do -- partly because she could **also** process the low level features of her own drawings in such a way as to compare them with low level percepts created by looking at the originals, also found difficult by most people who can see that their drawings are not very good but cannot fix them.

Learning to talk requires the ability to "chunk" sensory information to correspond to named objects properties, relations, and processes. Without such grouping processes the Popeye program could not have perceived lines, junctions, pairs of parallel lines, plates, plate-junctions and the groups of joined plates corresponding to capital letters. If teaching Nadia to talk had the effect of reallocating some of her information-processing resources to analysis and interpretation of more abstract features of her visual information, and setting up associations with sounds and muscular processes involved in producing words, that might have depleted the resources allocated to low-level vision, explaining why she lost some of her drawing skill. (Similar comments were made in a brief discussion of Nadia's unusual combination of abilities in Section 11 of [Sloman\(1989\)](#).)

Of course that is at best merely a plausible *sounding* explanation: it would have to be tested against both neural evidence, which may be hard to obtain, and against working models that demonstrate how vision and drawing capabilities might work. At present there is no robot that I know of that could draw as Nadia did, let alone one with the potential to be *taught* to speak as a child might be taught.

There is another information-processing distinction mentioned in [Sloman\(1989\)](#) that may be relevant to Nadia, namely the distinction between the "online" use of transient information to control action using visual feedback (visual servoing) and the storage of information for multiple potential uses later (e.g. in section 9.1). In general the information required for visual servoing will make use of "low-level" image details (e.g. the changing visible distance between, or the changing relative alignment of, two edges). So since Nadia could perform actions like picking up and manipulating a pencil -- which would have required such online control -- the resources might have been shared between that function and the functions involved in getting low level visual information required to support accurate drawing. Unusually powerful resources shared by mechanisms serving those two functions may have had something to do with her drawing competence, and her other limitations.

Note added 30 Mar 2013: Low level visual processing is not image construction

Some theorists may be tempted to interpret the suggestion that resources are allocated to low level visual information processing as implying that this is some sort of internal image construction. Nothing could be further from the truth, since a mere internal image could not provide the ability to create an accurate drawing. Images cannot control hand movements for example, or select a good sequence for drawing construction. The low level information processed may be *about* image structure

but that does not imply that it is *encoded in* some sort of image (as will be obvious to AI vision researchers, who know how much more is involved in image understanding than merely having an internal image.) [Max Clowes](#) once did an experiment with his children and some of their friends, the oldest aged 11 or 12 years, I think. He showed them pictures of impossible objects, like the Penrose triangle, and the Devil's pitchfork (often misleadingly labelled "illusions"), shown here (thanks to Wikimedia):

Figure Impossible (Added: 30 Mar 2013)

-----XX

Despite having the originals directly available, the children in the experiment, and some adults, all found the pictures very difficult to copy, without tracing. It would be very interesting to know how Nadia would have fared, or adults reading this. I once trained myself to draw the Penrose triangle by ignoring the 3-D structure and focusing on the 2-D structure: e.g. *there is a triangle at the centre with each of its sides extended in one direction, from which new lines can be drawn parallel to the previously existing sides, etc., etc.*

Using such a strategy may be much easier for someone not distracted by the 3-D interpretations of drawn fragments. **(To be expanded)**

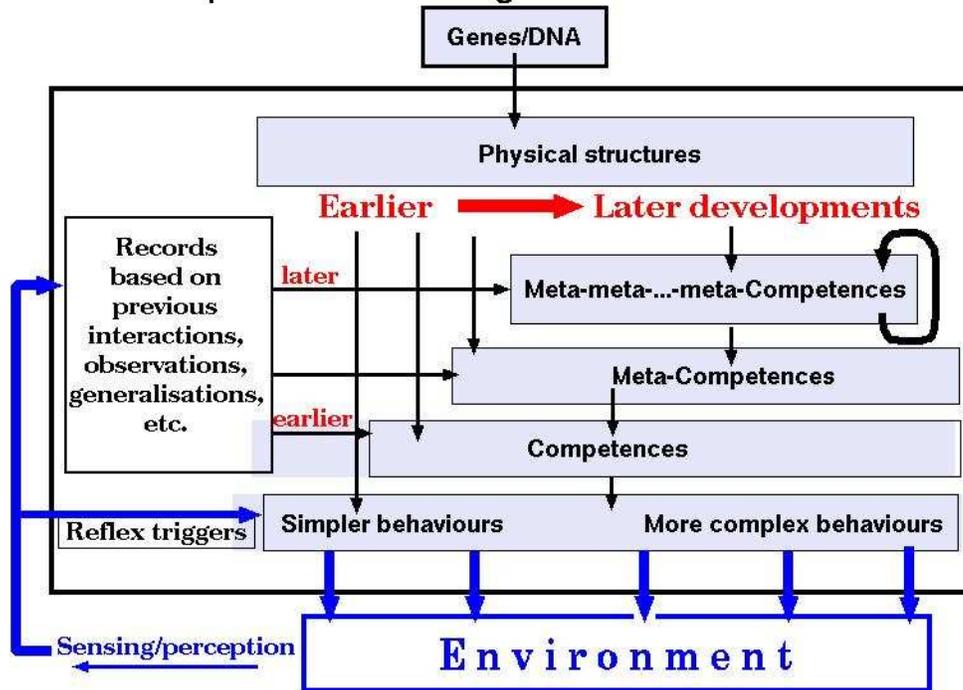
Developmental trajectories for "meta-configured" competences

There is a way of thinking about a wide range of developmental abnormalities that makes use of both the ideas about layered perceptual mechanisms illustrated by Popeye, albeit in a simple form, and more general ideas about layered competences developed thirty years later in collaboration with Jackie Chappell, [(2005), (2007)] to characterise the spectrum of differences between biological species described as "precocial" and those described as "altricial". In our [IJCAI 2005 paper](#) we suggested that instead of applying the distinction to species, or to members of species, we should apply them to competences: an individual could have a mixture of "precocial" competences mostly specified by the genome and available at birth or an appropriate developmental stage, and "altricial" competences that depend on high level specifications that come via the genome but are instantiated in a way that depends both on features of the environment and on previously acquired knowledge and competences. Later, in our [IJUC 2007 paper](#) we used the labels "pre-configured" and "meta-configured" to mark the difference.

Our sketchy theory is (crudely) summarised in the diagram **Fig EvoDevo** below, showing alternative routes from genome to behaviours, with increasing amounts, from left to right, of involvement of learning and development based on results of previous interaction with the environment, using previously developed competences and information stores (not shown). The pre-configured competences are on the left and increasingly meta-configured competences to the right, including possibly more complex routes further to the right, not shown in the diagram.

Figure EvoDevo

Multiple routes from genome to behaviours



[From Chappell & Sloman (2007)

which introduced the labels "pre-configured" and "meta-configured".

[Chris Miall helped with an early version of the diagram.]

This diagram is intended to apply not just to humans, but also to patterns of development in many other "altricial" species. Some individuals may go on developing longer than others, either because of genetic differences or because the environment continually re-stimulates certain learning/creating subsystems.

The diagram crudely and abstractly represents a huge space of possible patterns of learning and development. Over extended periods, different "parts" (not physical parts) of an individual may concurrently follow the loops round developing competences and then feeding back the results to influence the construction of more general and powerful types of competence or meta-competence. So processes of learning to walk, learning to see, learning to manipulate objects, learning to plan, learning to communicate with others, learning to make music, learning to do mathematics, learning to play various games, may develop in parallel, influencing one another in deep ways.

All species include "downward" developmental trajectories on the left, for some competences. Over time, in some species, evolution seems to have produced more and more of the routes further to the right, which depend on results of earlier learning based on competences developed via routes more to the left.

Routes on the left produce "pre-configured" competences. Routes further to the right produce competences that are "meta-configured, to greater extents.

However some sub-mechanisms cannot start developing until after others have amassed the information required for the new system to be useful. For example, an ability to switch from using learnt linguistic patterns to using more powerful generative syntax and semantic mechanisms should not start developing until a substantial amount of learning to use effective patterns of communication has provided enough data to drive a process of generative rule extraction without depending on

extrapolative guesses that later have to be rejected. Moreover after finding a powerful and apparently useful grammar, in a community of speakers who make mistakes, the process of accommodating exceptions to the grammar should not start until enough information has been collected to allow the exceptions to be distinguished from mistakes. Those two phases are part of the phenomenon known as "U-shaped" language learning, in which performance first develops to a high level then degrades, with new mistakes being generated, and then improves as the exception-handling mechanisms required to avoid those mistakes are developed (a non-trivial piece of software engineering for a young learner, or the learner's genome, to achieve!)

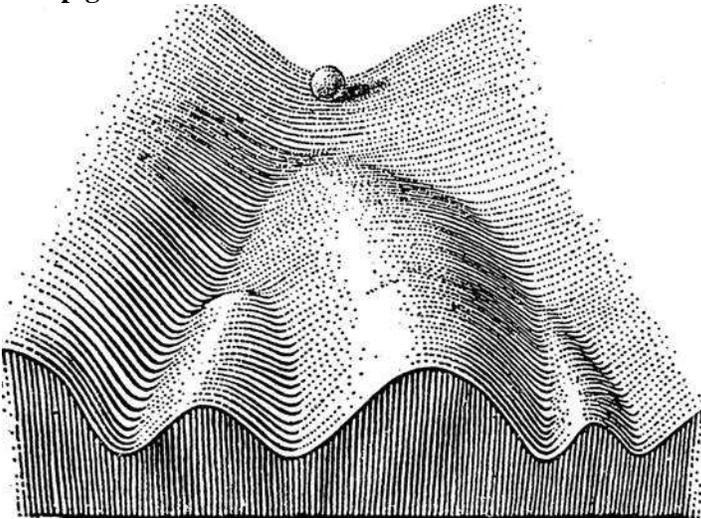
As Karmiloff-Smith notes, there's evidence that this pattern is not restricted to language-learning. There seem to be many forms of learning that require such layered learning, sometimes including construction of a temporary working system to provide a launch-pad for something more powerful. I suspect that close investigation of many animals with complex adult skills involving great variability in performance, such as hunting mammals, some nest-building birds, cetaceans, hunting mammals, elephants, and some foragers, will show that they do not move smoothly from infant incompetence to adult expertise, but develop skill and knowledge platforms that are required to support later developments. There is no assumption that uniform learning processes operate at all levels. On the contrary there seem to be enormous differences between different kinds of learning found in humans of different ages, from infants through toddlers to school and university learners to adult professionals of many sorts.

(Well-meaning educators who don't understand the information-processing requirements for learning mechanisms can sometimes introduce disastrous attempts to short cut development of "understanding" in a subject area, without first ensuring that a solid base of skill and knowledge has developed, including memorised facts and rules, which can be used to constrain and inform new pattern generation and hypothesis formation processes.)

Waddington's epigenetic landscape idea

The (recursive) map of possible alternative developmental trajectories can be seen as a much more flexible version of C..H. Waddington's notion of an "epigenetic landscape", often depicted like this:

Figure Epigen



[Figure by Waddington, displayed on many web pages without source]
[This version on <http://towardsdolly.wordpress.com/2012/08/>]

Here there's a fixed set of (steepest descent) routes from a starting position to lowest levels of the terrain (the end of life, perhaps). In contrast the trajectories in Fig EvoDevo produce behaviours that change future possible behaviours. And since the causal influences come not only from the genome but also repeatedly through the environment, changing future operations of the genome (e.g. late developing parts of the brain, or late developing kinds of virtual machinery) the space of possible developmental trajectories enabled by a particular genome can be huge.

That also implies that there is a huge space of possible deviations from biologically normal functioning, whether caused by a chemical or other minor abnormality at an early stage with consequences fanning out later, or caused by something in the environment that can distort future branches in the space of possible developments.

An implication (that not everyone will find obvious) is that characterising those abnormalities simply in terms of their physical and behavioural manifestations will not give us a deep understanding of what they are and what their later implications are. Instead we need to base the characterisation of particular abnormalities on both their epigenetic history (which can be mixture of nature and nurture) and their possible epigenetic futures. Compare characterising a sentence by both its parse tree and the implications that can be drawn from it.

The need for multi-level architectures

Another fact that adds to the complexity of the mechanisms and processes described here is the concurrent operation of different kinds of information processing in perceptual systems, action systems, and more central systems, corresponding to brain functions that seem to have evolved at different times.

Many researchers in AI/Cognitive science have over several decades, suggested that robots and animals need information processing architectures in which different layers of competence are used in parallel. Some researchers attempt to design the architectures on the basis of analysis of task requirements, some just follow fashions and available toolkits, and some explicitly attempt to build architectures based on evidence or theories about natural cognitive systems. Some of the latter group were brought together a few years ago by the BICA project (Biologically Inspired Cognitive Architectures). See

- <http://bicasociety.org/> Society web page
- <http://bicasociety.org/cogarch/> Online table showing a variety of biologically inspired architectures.

Unfortunately there are fads, prejudices and constraints related to available tools, so that very few researchers are aware of what all the others others have tried and why.

The Birmingham Cognition and Affect (CogAff) [project](#) developed a theory emphasising three main sub-divisions of information-processing mechanisms of varying age, though each could be subdivided into more specialised layers or sub-systems and it may be necessary to add more types of mechanism. The oldest layer also found in insects and other invertebrates is sometimes called the "Reactive" layer, which is incapable of representing things that do not exist, e.g. alternative possibilities or possible future actions and their consequences. The "Deliberative" layer includes mechanisms with those capabilities missing from reactive layers, and uses them in many different ways, including constructing multi-step plans. The "Meta-Management" (meta-semantic) layer, which is assumed to be the most recently evolved sub-system can represent and reason about things that represent and reason including oneself and others ([Beaudoin\(1994\)](#)).

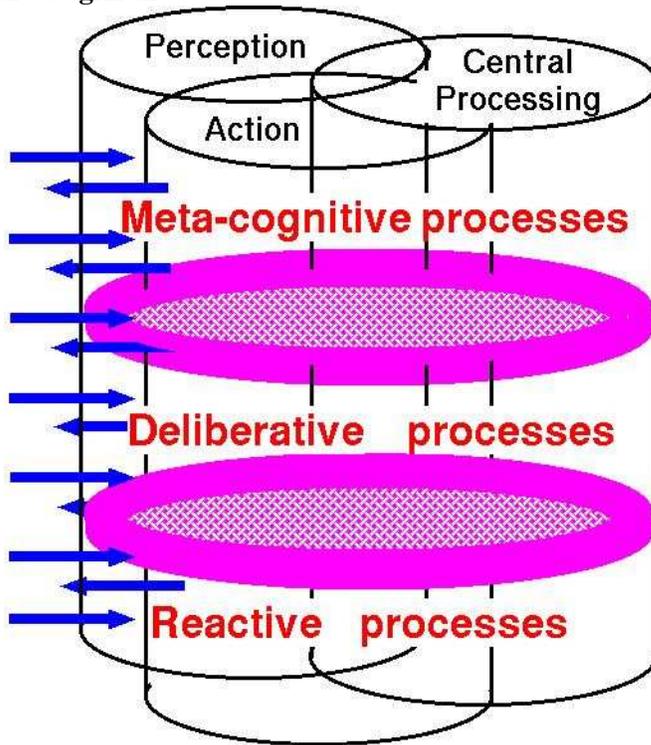
Note on Concurrency added 8 Apr 2013

N.B. this scheme is not intended to imply that there is some kind of pipeline of processing of different types from sensors through to motors, as is sometimes implied by AI theories emphasising a "Sense, Think, Act" loop. The CogAff schema includes architectures in which the different layers operate in parallel (even if they do not all develop in parallel). For example, it is important for some parts of the system to monitor what is currently happening in other parts, instead of only examining traces of activity after the event.

Moreover, even within a layer, there will usually many different activities going on in parallel, e.g. perceiving various changes in the environment while thinking about future trajectories and while controlling current walking or running. In addition, "Alarm" mechanisms may be able to monitor aspects of activity in all the different layers and in some cases to take control and redirect processing. This idea is part of a theory of varieties of emotional and other affective processes.

That layering of functions would be manifested not only in divisions within central processing mechanisms but also within perception and action subsystems, leading to a class of architectures that could be (crudely) represented as in **Fig CogArch** below, with multiple, input and output streams operating in parallel.

Figure CogArch



(With thanks to Dean Petters.)

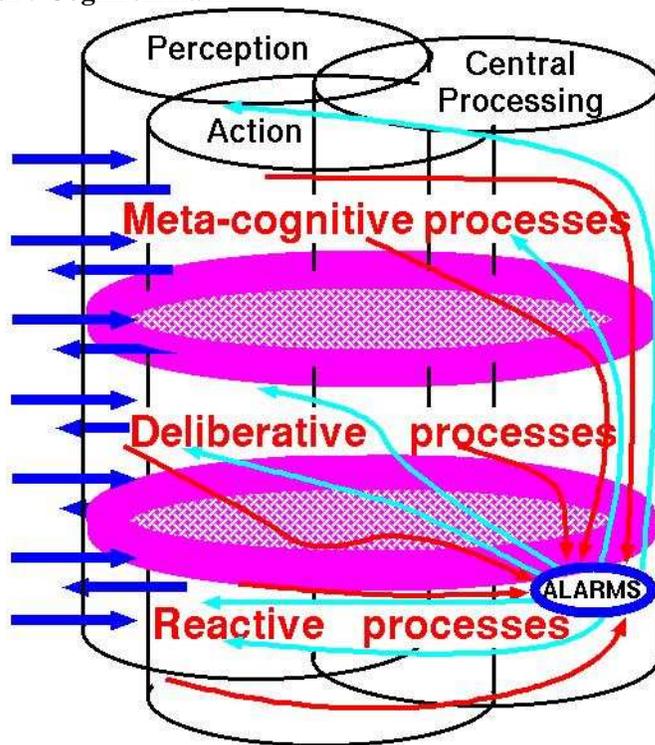
If there are (at least) three very different kinds of information-processing subsystems with different evolutionary histories and those differences are also found in perception, action and more central subsystems, and all the subsystems share resources, as indicated by the overlap in **Fig CogArch** then the scope for variable patterns of development is enormous, as all those systems grow and mature in parallel. That's an understatement!

Added 8 Apr 2013: Alarm mechanisms and varieties of emotion

A further complication, developed in papers and presentations in the CogAff project, is that there can be "Alarm" mechanisms of various kinds, capable of receiving information from all parts of the system, assessing it quickly, and in some cases taking control and rapidly redirecting some or all of the other subsystems, with results that might include fighting, feeding, fleeing, freezing or sexual activity. (Often "freezing" is omitted from the so-called "Four Fs".) The alarm mechanisms are essentially reactive mechanisms, because there is often no time for deliberation and higher level evaluation and decision-making processes. This implies that the actions of such a system are inherently unintelligent and may sometimes lead to disasters, as pointed out in ([Sloman and Croucher, 1981](#))

[Figure CogArchAlarm](#), below is a crude representation of possible interaction routes involving Alarm mechanisms, superimposed on [Figure CogArch](#), above. Depending on which architectural layers are involved these mechanisms can produce primary, secondary or tertiary types of emotions. Moreover, the states are inherently **dispositional**. That is, in some cases instead of the Alarm mechanisms actually taking control, they have a disposition or tendency to take control, but that may be overridden by some other mechanism, possibly driven by a very high priority motive, and making use of acquired meta-management control capabilities.

Figure CogArchAlarm



It is very likely that there are types of abnormality that result from interactions between these alarm mechanisms and other mechanisms, including cases where one sub-system has developed "normally" but not others. This could cause considerable control difficulties for carers.

Requirements for explanations of cognitive abnormalities

The only way to gain a deep scientific understanding of autism, and other developmental abnormalities affecting mental functioning, including Down Syndrome, Williams Syndrome, Tourette's syndrome, and probably many others, is to develop a deep theory of varieties of biological information-processing, within which it is possible to explain how human information-processing mechanisms and competences develop in individuals, from fertilization onwards. On that basis, we can begin to ask and answer clear questions about **variations** in developmental trajectories, and their implications. Finding causes and removing them, or reducing their effects, may be much harder: but understanding what's actually going on, as opposed to what family members, teachers, the individuals concerned, and others *experience* as going on, is a pre-requisite for any informed effective policy, whether it's a policy for prevention, for amelioration, or for coping. It is also a requirement for science, as opposed to sympathy.

Often people ask why something is not the case, when they would not ask why it is the case otherwise. For example, people may ask why a particular child does not grow normally, but they don't ask why other children do grow normally. When shown examples of change-blindness (as described in http://en.wikipedia.org/wiki/Change_blindness) they tend to ask "Why can't we see the changes that occur?" when one picture is replaced by another but they don't ask in more normal situations "How can we see changes when they occur?". Seeing a shadow appear and reappear is the normal case, so we don't ask why it happens.

But change perception requires more mechanisms than perception of a scene at a moment. It requires the ability to compare the current state with a previous state, including, in some cases, the ability to search for a small difference between the two. Without understanding the normal mechanisms and why they don't work as normal in the change blindness experiments we cannot explain change blindness.

Likewise, without understanding the mechanisms and processes involved in normal cognitive abilities and their development we cannot hope to explain the abnormal cases. Even if it were true that a particular type of diet produced some variety of autism, that would not explain what the differences are between individuals with that sort of autism and normal individuals, nor why the diet produces those differences.

Further, I suggest that it is impossible to gain a deep understanding of the abnormal phenomena without having an equally deep understanding of the phenomena they are compared with. For example, trying to understand what's wrong with a child who cannot interact socially with others, or who cannot learn a human language, or whose grasp of spatial reasoning is inadequate for a normal life, without relating the answers to the mechanisms that make it possible for others to interact socially, learn a language, grasp spatial reasoning, can be compared with trying to understand why some aeroplane engines fail in cold weather without understanding how the other engines work in all weathers. That understanding will not come merely by observing whole functioning engines in various conditions: what's going on inside that enables them to perform their functions is crucial. Similarly what's going on inside a language learner or a social agent, or a spatial reasoner is crucial to understand how things can go wrong.

Unfortunately, until very recently, the available ways of thinking about what's going on inside a very complex information-processing machine were grossly inadequate for the purpose. Even now, although much has been learnt about how to design, develop, and debug, artificial information processing systems, most of the professionals concerned with studying or caring for humans remain entirely ignorant of what has been learnt, and the people who have learnt the new concepts and

theories are more interested in using them to build new useful machines than to explain the information processing products of biological evolution and development: a much harder, and (at present) less financially rewarding, task.

The need for a Meta-Morphogenesis project

Since humans are products of millions of years of evolution, producing millions of "design decisions", what exists today is hugely complex and hard to investigate or to understand. I have proposed the [Meta-Morphogenesis project](#) (inspired by Turing's 1952 paper on morphogenesis) as a means of taming the complexity, on the basis of descriptions and explanations of changes in biological information processing since the dawn of evolution on this planet. If we can learn more about the design changes required to transform much simpler types of information processing into the kinds that now exist, that may help us to understand some of the hidden complexity in the mechanisms that now work together to enable and control process of learning, development, perception, motive formation, action selection, control of actions, theory formation, and many more.

That is as true for problems connected with mental/cognitive functioning as for problems connected with physiological functioning, such as diabetes, congenital heart deformities, allergies, or asthma, although it is sometimes possible to discover an ameliorative treatment for a problem without understanding the nature of the problem or why the treatment works, such as discovering that a particular chemical substance reduces the pain of some headaches without knowing what the headaches really are or how the chemical works. (In other words, "craft" knowledge can precede scientific understanding. But scientific understanding can be the basis of "engineered" solutions, that are both more effective, better targeted, and better understood than craft solutions.)

I do not mean to disparage detailed reports by parents, teachers, doctors, social workers, siblings, friends, clinical psychologists, or informed novelists. [Compare Hacking\(2009\)](#) I merely wish to emphasize the theoretical and practical importance of what such studies leave out, like reports of cancer before the development of modern biochemical epigenetic research.

It's hard partly because, unlike morphology and behaviours, which many others have studied, forms of information processing, and the mechanisms required, are very hard to observe, even when the individuals concerned are alive and functioning now. That's partly because I-P machinery has become increasingly dependent on virtual machinery implemented via layers of older virtual machinery. So neither observing external behaviours nor observing internal brain processes can reveal what's going on.

The project therefore has to be highly conjectural, but constrained by as much evidence as possible (e.g. how giraffes use their tongues to get leaves from acacia trees full of potentially dangerous thorns, how some ants teach others the route to a new nest, how squirrels manage to defeat squirrel-proof bird feeders, etc., how information processing requirements for increasingly sophisticated genomes have changed, etc.).

I think the project can be constrained so as to constitute a "progressive" research programme, in the sense of [Lakatos \(1980\)](#)

The M-M project has spawned a number of strands, including stuff about emergence of euclidean geometry, toddler theorems, development of virtual machinery, changing mechanisms of motivation, changing information-processing architectures, changing forms of representation, changing ontologies used, ... and there will have to be a growing network of growing web pages. Perhaps the project will later spawn a much deeper theory about varieties of Autism and many other products of diverse developmental trajectories.

(Added: 29 Mar 2013) Notes on Motivation (incomplete)

Earlier I mentioned mechanisms of motivation briefly. Although it is possible to study visual and other forms of information-processing in organisms without being concerned with motivations driving or served by those processes, a common assumption is that whatever internal or external actions are performed by an organism must be chosen because of some expected reward, which could be either be something wanted for itself or some form of reduction or avoidance of something not wanted, e.g. pain, hunger, displeasure.

I think the assumption that there's no motivation without expectation of reward is false, and that evolutionary considerations as well as close observation of children and other animals suggest that there is another kind of motivation at work, which is motivation generated by a mechanism chosen by evolution because the effects are found to be useful, but without the individual having any knowledge of those mechanisms, or what they do or why.

I have tried to develop this idea in [Sloman\(2009\)](#), though many readers confuse this kind of motivation with motivations related to genetically determined rewards, e.g. satisfying curiosity, having fun, or whatever. That's because the phenomenon of doing something just because you want to do *that*, and not because doing that will produce *something else*, tends to go unnoticed, even though, as Gilbert Ryle pointed out in *The Concept of Mind*(1949), it is a common feature of hobbies, extreme sports, music making, doing mathematics, making things, and climbing mountains. One consequence of young animals having such architecture-based motives (generated by mechanisms deep in their information-processing architecture) is that they do a great deal of learning that they would not otherwise do, much of which is useful later. But they cannot possibly know this at time of acting, so that benefit cannot their reason for acting, even if that was the cause of evolution's selecting the mechanism: individuals with it tended to have more offspring than those without (other things being equal).

From the reports on Nadia I have read it appears that her motivation to make drawings was of that architecture-based kind, and attempts to interpret the behaviour as designed to achieve something else are generally based on the prejudice that wanting merely to do what she does, for no ulterior reason, is somehow impossible.

I've recently encountered a researcher, Emre Ugur, who came up with the idea I have labelled "architecture-based motivation" independently and moreover has gone much further than I have insofar as he implemented a working robot with such architecture-based motives, which he unfortunately labelled "intrinsic motives", a label often used by proponents of a kind of reward that is innately desired, for which various actions are means. His work, demonstrating how possession of such motive generators can lead to important forms of learning is presented in his PhD thesis [Ugur\(2010\)](#) and subsequent publications.

(To be expanded/revised.)

[Back to CONTENTS](#)

Note added: 2 Jan 2014

I have just encountered this very interesting and relevant online discussion:

"Are Prodigies Autistic?" by Scott Barry Kaufman

<http://www.psychologytoday.com/blog/beautiful-minds/201207/are-prodigies-autistic>

Here's an extract:

"More striking is that every single prodigy scored off the charts in working memory -- better than 99 percent of the general population. In fact, six out of the eight prodigies scored at the 99.9th percentile! Working memory isn't solely the ability to memorize a string of digits. That's short-term memory. Instead, working memory involves the ability to hold information in memory while being able to manipulate and process other incoming information. On the Stanford-Binet IQ test, working memory is measured in both the verbal and non-verbal domains and includes tasks such as processing sentences while having to remember the last word of each sentence, and recalling the location of blocks and numbers in the correct order in which they were presented. There have been many descriptions of the phenomenal working memory of prodigies, including a historical description of Mozart that involves his superior ability to memorize musical pieces and manipulate scores in his head."

(To be expanded/)

[Back to CONTENTS](#)

Some references (incomplete)

(Later to be divided into (a) closely connected with this paper (b) background.)

- Ian Apperly, **Mindreaders: The Cognitive Basis of Theory of Mind**, Psychology Press, London, 2010,
- Noam Chomsky, *Aspects of the theory of syntax*, MIT Press, Cambridge, MA, 1965,
- <http://www.ncbi.nlm.nih.gov/pubmed/19528032>
Ian Hacking, Autistic autobiography.
In *Philosophical Transactions of the Royal Society, Biological Sciences*, 364 (No. 1522 27 May) (2009): 1467-1473.
Reprinted in R. Happi and U. Frith, *Autism and Talent*, OUP, 2009, pp. 195-207.
- Aaron Sloman
The Computer Revolution in Philosophy: Philosophy, science and models of mind
Harvester Press and Humanities Press, 1978
now online with notes and comments added intermittently after 2001, here:
<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/>
- <http://www.cs.bham.ac.uk/research/cogaff/81-95.html#36>
Aaron Sloman and Monica Croucher, "Why robots will have emotions",
Proc 7th Int. Joint Conference on AI (IJCAI-81) pp. 197--202,
Vancouver, 1981
(Some readers insist on interpreting this paper as recommending that robots **should** have emotions!)

- <http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#7>
A. Sloman, On designing a visual system
(Towards a Gibsonian computational model of vision),
Journal of Experimental and Theoretical AI, 1, 4, 1989, pp. 289--337,
 - <http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#907>
Aaron Sloman,
Architecture-Based Motivation vs Reward-Based Motivation,
in *Newsletter on Philosophy and Computers*,
American Philosophical Association, 09, 1, 2009, pp. 10--13,
 - <http://www.cs.bham.ac.uk/research/projects/cogaff/#overview>
A. Sloman and colleagues: Origins and overview of the CogAff project (1991--)
 - <http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>
Aaron Sloman and Jackie Chappell, The Altricial-Precocial Spectrum for Robots,
in *Proceedings IJCAI'05*, Edinburgh 2005, pp. 1187--1192,
 - <http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#717>
Jackie Chappell and Aaron Sloman,
Natural and artificial meta-configured altricial information-processing systems,
International Journal of Unconventional Computing, 3, 3, 2007, pp. 211--239,
 - <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703>
Aaron Sloman and Jackie Chappell, Computational Cognitive Epigenetics
(Commentary on Jablonka and Lamb), in *Behavioral and Brain Sciences*,
30, 4, 2007, pp. 375--6,
 - <http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#38>
Luc P. Beaudoin, *Goal processing in autonomous agents*,
PhD Thesis, School of Computer Science, Birmingham, UK, 1994,
 - C. H. Waddington, **The Strategy of the Genes**, 1957, MacMillan,
 - Eva Jablonka and Marion J. Lamb,
Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life,
MIT Press, Cambridge MA, 2005,
 - Imre Lakatos, Falsification and the methodology of scientific research programmes,
in *Philosophical papers, Vol I*, Eds. J. Worrall and G. Currie,
Cambridge University Press, 1980, pp. 8--101,
 - <http://www.cns.atr.jp/~emre/papers/PhDThesis.pdf>
Emre Ugur,
A Developmental Framework for Learning Affordances
PhD Thesis The Graduate School of Natural and Applied Sciences
Middle East Technical University, Ankara, Turkey, 2010,
-

Related items

- <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/alchemy/>
Is education research a form of alchemy?
Originally posted 08 June 2012. [Issue 27 of ALT Newsletter](#)
(On the risks of expecting too much from theoretically impoverished empirical research.)
- <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>
An introduction to "Virtual Machine Functionalism" as a relatively new conceptual framework for thinking about old problems about minds, learning, development, and evolution.
- Three video presentations that may help to explain the background to some of the thinking here
 - <http://www.youtube.com/watch?v=QXAFz3L2Qpo>
What is computational thinking? Who needs it? Why? How can it be learnt?"
Invited talk at the 2012 conference of the Association for Learning Technology (ALT), Manchester, 11 September 2012.
 - <http://www.youtube.com/watch?v=5dEXIOiAsaw>
Adam Ford interviews Margaret Boden at the AGI 2012 Conference, Oxford December 2012.
How AI can inspire new better questions about minds, compared with older approaches.
Talks about Autism and theory of mind -- briefly mentions work of Chris Frith on differences in processing psychological verbs in autists and others.
 - <http://www.youtube.com/watch?v=iuH8dC7Sjno>
Adam Ford interviews Aaron Sloman at the AGI 2012 Conference, Oxford December 2012.
A more general introduction to ways of thinking about information processing and biological evolution.

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