
THIS WILL BE MODIFIED FROM TIME TO TIME

2015 onwards:

**Afterthoughts (HTML and PDF) on
THE COMPUTER REVOLUTION IN PHILOSOPHY
Philosophy, science and models of mind.**

The book was originally published in 1978.

Free online revised version available since July 2015 (HTML and PDF)

(Some notes, references and re-formatting added since then.)

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

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

Afterthoughts relating to the 1978 book will be added to this document (started July 2015) possibly over the next few years, including summaries of and references to later developments of some of the ideas.

Aaron Sloman

The School of Computer Science

University of Birmingham, UK

CONTENTS

- [Major development, related to hatching processes in eggs \(2021\)](#)
 - [1978 -- 2015](#)
 - [Origins of this new, free, online edition](#)
 - [Afterthoughts](#)
 - [A key \(Kantian\) idea: Explaining Possibilities \(Chapter two\).](#)
 - [Possibilities \(tentatively and sketchily\) explained in the book](#)
 - [The Meta-Morphogenesis Project \(originally proposed in 2012\)](#)
 - [Themes from the 1978 book developed in subsequent work](#)
 - [References](#)
-

Major development, related to hatching processes in eggs (2021)

For details see this paper, prepared for a talk in September 2021:

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

NOTES FOR PRESENTATION AT MORCOM Sept 15 2021

Why Don't Hatching Alligator Eggs Ever Produce Chicks?

Animal spatial intelligence and limits of current AI and neuroscience.

1978 -- 2015

The Computer Revolution in Philosophy (CRP) was published at the instigation of colleague and friend Margaret Boden, by Harvester Press (UK) and Humanities Press (USA) in 1978, when the author was at the university of Sussex. It has been out of print for some time, though used copies are still offered for sale.

Thanks to much appreciated labours of colleagues as described below, versions of most parts of the book have been available online since 2001, initially as separate HTML files (with many scanning and OCR errors that were gradually corrected), and later converted to PDF. A few notes and comments were added from time to time.

During 2015 work began on combining all the chapters into a single document with internal cross references, with some minor new corrections and a few comments that had previously been inserted in the electronic versions of the chapters. The new integrated version of the book, including an integrated table of contents, is available in two formats, HTML and PDF, made freely available in August 2015, with a creative commons licence:

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

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

The table of contents of the integrated online version, included in both formats, can be viewed here:

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

Note:

An early online version of this book did not include the original index (pages 288-304). A searchable version of the index was added in 2016. It is also available as a separate collection of images available in two formats:

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/index-pages.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/index-pages.pdf>

The need for an index is reduced by the fact that the electronic versions can easily be searched for key words or phrases.

There are several other online versions of the book, that are now all out of date when I last looked at them. They had all been copied (with or without explicit permission) from the freely available "master" version in various formats since 2001, including:

the multi format version at 'archive.org':

<https://archive.org/details/TheComputerRevolutionInPhilosophyPhilosophyScienceAndModelsOfMind>

the Amazon Kindle version kindly produced by Sergei Kaunov (with a very kind review) in 2011:

<http://www.amazon.co.uk/The-Computer-Revolution-Philosophy-science-ebook/dp/B006JT8FSK>

and others referenced here:

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

Origins of this new, free, online edition

This new version of *The Computer Revolution in Philosophy* was made possible because in 2001 Manuela Viezzer, a PhD student in Cognitive Science, scanned my last copy of the paper version, and Sammy Snow, a departmental administrator, used OCR software on her PC to produce a first RTF version, which I gradually proofread, edited, converted to html, corrected in some places, and supplemented with additional notes and references, all made freely available as HTML and PDF, with a Creative Commons licence. For several years, the preface, chapters, epilogue and postscript were in separate files, in HTML and PDF formats, though a combined PDF file was later made available for download.

Note: Manuela Viezzer is now an artist:

<https://manuela-viezzer.squarespace.com/>

<http://www.saatchiart.com/manuela.viezzer>

From time to time various corrections and additions were made. The additions were indicated in the text, though minor corrections were not all labelled. In some cases the additions were notes at ends of chapters, though a few were inserted (usually clearly indicated) in the main text, adding to or (hopefully) clarifying what was said.

Fourteen years later, all the various parts except for the name/subject index (which is still in scanned image mode, not text mode) were combined to form a single html file, along with a derived pdf file -- thanks to two excellent (free) linux programs, html2ps and ps2pdf.

Afterthoughts

This "Afterthoughts" document, begun 22 Jul 2015, and likely to be extended from time to time, is available in two formats, HTML and PDF (derived from the HTML version):

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/crp-afterthoughts.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/crp-afterthoughts.pdf>

The 1978 book had many flaws of presentation, some of them formatting flaws (e.g. arbitrary changes of font or indentation), and others to do with the fact that the ideas needed far more development. I hope the new version improves on the format at least. This document attempts to improve on some of the flaws of presentation of ideas in the original book, and to summarise some of the work that grew out of the book over the next few decades, with pointers to relevant publications.

One of the features of the book that seemed to annoy some readers and reviewers was that it made no attempt to survey the achievements of Artificial Intelligence, and instead focused on gaps and unsolved problems, along with suggestions for making progress, often requiring long term research. Perhaps I was foolish, but I assumed that excellent introductory overviews and edited collections of papers, produced by others especially Margaret Boden (*Artificial Intelligence and Natural Man*, also published in 1978), and collections edited by Feigenbaum and Feldman, Minsky, McCarthy and others provided adequate introductions to the rapidly expanding field.

Perhaps the worst fault was the arrogance and intolerance displayed in several places, where I reacted to theories I felt were shallow, arguments that I thought were inconclusive, and prejudices against new forms of explanation, including prejudices against computational answers to

philosophical questions and explanations making use of new ideas about virtual machinery.

I was also strongly opposed to a very popular view of science (which is still too popular) deeply influenced by Popper's falsificationism, which I felt did not do justice to key features of many great scientific advances. That was the belief that only empirically falsifiable statements could be part of scientific knowledge. Popper himself eventually recognized flaws in the requirement that scientific theories be empirically falsifiable, for example when he recognized the great scientific merit of Darwin's theory of natural selection, and even began to speculate about evolutionary mechanisms himself, e.g. in Popper (1978).

Chapter 2 of the book ("What are the aims of science?") attempted to present an alternative to Popper's falsificationism, but apparently failed to communicate with most readers. More on that below.

Despite the arrogance and intolerance of the style, and the apparent difficulty most readers had in appreciating Chapter 2, on which everything else rested, some readers, including the reviewers referenced below, recognized that important new ideas were being discussed in the book. Those ideas have been developed further in other publications since 1978, but some students and researchers may find it useful to have the earliest versions of the ideas readily available in this free edition.

Not all the hopes at the time of writing have been justified. In particular, there remain important features of animal perception that are closely related to human abilities to make mathematical discoveries, that, as far as I know, have not yet been modelled in AI systems and seem to resist such modelling for reasons that are not entirely clear. For example, Chapter 7 of CRP attempted to demonstrate the importance of non-Fregean (i.e. "analogical") forms of representation in human intelligence, but good working implementations of the ideas proposed do not yet exist (as far as I know). As a result we cannot yet model the processes of discovery that originally led to Euclid's *Elements* about 2,500 years ago. Some of the required visual mechanisms are illustrated and discussed in:

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

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

Some (Possibly) New Considerations Regarding Impossible Objects

Their significance for mathematical cognition,

and current serious limitations of AI vision systems.

(References to further work in progress on that problem will be added here later.)

Hard to model mental causation

A more subtle problem that I believe remains unsolved is the gap between forms of causation in human and animal minds and the attempts to model them computationally. A stark example is the causal state in which someone has two strong, but opposed desires, e.g. a desire to take revenge on someone and a desire to avoid the consequences of doing so (e.g. losing the affection of the victim's sibling), or the desire to eat a very tempting desert and the desire to lose weight. In computer models such conflicts are typically represented by using numbers to represent strengths of the desires, and letting the stronger desire (i.e. the one with the higher numerical measure) dominate the weaker desire. Once the comparison has been made the stronger desire is allowed to "win" and the corresponding actions are executed. The weaker desire (typically) then plays no further causal role.

But in humans, and presumably some other animals, the fact of choosing one desire does not make other desires inactive. A "rejected" desire can remain as strong as it was despite the other desire having been selected for action, and the rejected desire can go on interfering with reasoning and actions. It is fairly obvious how simple forms of persistent conflicts might be implemented in neural nets, but it is not clear how the richness of human experiences of conflict and indecision can be explained, nor what happens to decisively rejected desires. An incomplete discussion of these ideas can be found in a draft online presentation here (work in progress):

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk86>

Supervenience and Causation in Virtual Machinery

A key (Kantian) idea: Explaining Possibilities (Chapter two).

Since 1978 the world has moved on and I've learnt a great deal more than I knew while writing the book. But many of the key ideas of the book, including ideas criticised by reviewers, still seem to me to be important even though they were explicitly rejected by some readers, and simply ignored by many others. Of these, one of the most important themes, the key claim of Chapter 2 of CRP, is that discovering and explaining *possibilities* is a more basic function of science, and more influential in the long term, than discovering and explaining laws or regularities, including statistical regularities. This feature of science implies and explains deep overlaps between science and philosophy.

I have therefore included, alongside the new online edition of the book, a paper written in part in response to the criticisms of Chapter 2 in reviews by Steven Stich and Douglas Hofstadter (both of whom apparently liked other aspects of the book). The new paper (new in November 2014) can be found here:

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

Construction kits as explanations of possibilities For links to the reviews see

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/crp-reviews.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/crp-reviews.pdf>

As far as I know only two people ever agreed with the main claims of Chapter Two: Trevor Pateman, a fellow philosopher at Sussex who was one of the founding editors of the *Radical Philosophy Journal* (which published an earlier draft of the chapter) and Tony Leggett, a philosophy (Lit. Hum. Oxford) graduate who later became a distinguished theoretical physicist. He alludes briefly to our interactions in this autobiographical note:

http://history.phys.susx.ac.uk/Research-Tony_Leggett and in the Preface to his 1987 book. At the time I found his approval very encouraging.

NOTE ADDED 19 Oct 2015

A tentative discussion paper shows a connection between having an explanation of a collection of related possibilities and a useful strategy for assessing competences based on those possibilities.

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

Possibilities (tentatively and sketchily) explained in the book

Chapter 2 of CRP -- now freely available online as part of the new electronic edition -- claimed that explaining how something (or some class of things) is possible is a major function of science. Many major past scientific advances were theories about what is possible and explanations of some possibilities in terms of more fundamental possibilities.

After further methodological preliminaries, the rest of the book presented (sometimes tentative) examples illustrating how AI (including computational linguistics) could advance our ability to explain (and sometimes predict) possibilities, as theories in physics and chemistry had done previously. I deliberately chose phenomena that had not yet been satisfactorily explained to show how developments in AI might advance understanding, rather than presenting examples of past achievements in AI, as several other authors had done.

For example, Chapter 6 attempted to explain (only in outline) how a machine with a certain sort of mind (more specifically: with a certain sort of information-processing architecture, with changing, interacting, concurrently active, components, directly or indirectly linked to sensors and effectors) could deal creatively with a complex, changing and locally unpredictable universe. Those architectural ideas were later elaborated with the help of a succession of PhD students and colleagues, after I moved to Birmingham in 1991, and started the "Cognition and Affect (CogAff)" project, initially called "Attention and Affect" and based on collaboration with Glyn Humphreys (then head of psychology in Birmingham). Some of the ideas in the CogAff project are summarised in another document, here: <http://www.cs.bham.ac.uk/research/projects/cogaff/#overview> and further developed in later presentations and papers, including an overview of Virtual Machine Functionalism (VMF).

Chapter 6 distinguished two high level "loops", which could run in parallel or alternating, labelled "the executive loop", concerned with carrying out detailed actions that had been selected, and "the deliberative loop", which was concerned with considering alternative options in response to a variety of types of events (including interrupts, goals being achieved, failures being detected, new options detected, etc.) that could indicate that something in the executive loop needed to be modified, temporarily interrupted, terminated or re-directed. I later learnt that the ideas in this chapter (circulated before publication) had influenced work by Tim Shallice.

However, most psychologists and neuroscientists paid no attention, though I later learnt that cognitive psychologists and various researchers on mental disorders had chosen the label "executive" for the mechanisms and processes that I called "deliberative".

Chapter 7 attempted to explain how valid reasoning, using physical mechanisms or mental (virtual machine) systems could validly and fruitfully use non-Fregean forms of representation including both maps and diagrams depicting physical machines. Fregean representations are composed entirely of functions (including higher order functions) applied to arguments (which could also be functions). I call such representations "Fregean" because Gottlob Frege was the first person (as far as I know) to identify the full generality of the concept of a function, including showing predicates and relation words in ordinary language can be treated as functions whose results are truth values (i.e. true or false) and how the universal and existential quantifiers ("all" and "some", or "there exists") can be interpreted as higher order functions applied to predicates and relations. (My 1962 DPhil thesis generalised this to allow "the state of the world" to be an additional implicit argument.)

Chapter 7 was originally a slightly modified version of a paper on Fregean vs Analogical representations presented in 1971 at the 2nd International Joint Conference on Artificial Intelligence (IJCAI), and published later that year in the journal *Artificial Intelligence* vol 2, 3-4, pp 209-225, 1971). The paper was primarily a criticism of the claim by McCarthy and Hayes (1969) that a notation based on first order logic would be adequate for an intelligent robot. (They distinguished three levels of adequacy, Metaphysical, Epistemological and Heuristic.) My paper (and the version in Chapter 7) aimed to show that whereas the logical notations were Fregean (i.e. using function/argument relationships), in some cases a different sort of notation, which I called

"analogical" where properties and relationships represented properties and relationships, possibly in a context-sensitive manner, would have advantages, especially heuristic advantages.

Many other thinkers made related distinctions before that and after that, though in many cases the subdivision was mistakenly described as involving continuity vs discontinuity, or a vaguely defined distinction between *symbolising* and *being similar*.

Often the role of analogical representations was confused with use of isomorphism or similarity between representation and things represented (ignoring all the dissimilarities between 2-D images and 3-D structures they represent, discussed in the chapter). The role of non-Fregean representations in mathematical discovery and reasoning remains largely unexplained, and human performances are still not even closely approximated by AI systems developed so far. Several examples are presented in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/impossible.html> and related documents.

The themes in Chapters 7 and 8 are related to work begun two decades earlier on my Oxford DPhil thesis, completed in 1962:

[Knowing and Understanding:](#)

[Relations between meaning and truth, meaning and necessary truth](#)

I was too academically naive to realise that the thesis should be published in book form. However, thanks to much help from Luc Beaudoin, it was eventually digitised with searchable text in 2016, and is now available online here in plain text and PDF formats:

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

That work is still in progress. E.g. the key questions about the nature of still unexplained mathematical discoveries made centuries ago by Euclid, Archimedes and many others are discussed in this conference presentation.

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.html>

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/ijcai-2017-cog.pdf>

I learnt in 2018 that Alan Turing had made related points in his thesis (published in 1938) where he distinguished mathematical *intuition* and mathematical *ingenuity* and claimed that computers (e.g. Turing machines) were capable only of mathematical ingenuity, not mathematical intuition, though he did not say why. His claim is discussed in

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/turing-intuition.html> (also PDF).

Chapter 8 attempted to explain (in very sketchy outline) how a young learner could begin to learn about numbers, by learning about one-one correspondences -- a dependence pointed out by David Hume, and exploited in the great work of Frege, Russell and others attempting to reduce arithmetic to logic -- completely ignoring the mental mechanisms required for uses of number concepts, a gap the chapter aimed to fill, at least in outline, since many details were not mentioned.

The chapter explained various conjectured procedural and representational mechanisms available for using numbers in counting and reasoning activities, including mechanisms for concurrent execution of procedures (e.g. pointing at objects while reciting number names), and mechanisms for observing and controlling concurrently running sub-systems. It attempted to show how a child (for example) could begin to use a memorised sequence of arbitrary symbols, to perform a variety of tasks involving setting up one-to-one correspondences over time, or correspondences in spatial configurations, and then go on later to discover "theorems" about numbers and new procedures for using them.

Unlike many psychological investigations of uses of number concepts this chapter emphasised the central importance of one-one correlations, and the variety of mechanisms for detecting and using them, mostly ignored by psychologists and neuroscientists as far as I know. One researcher who seems to have independently developed a closely related approach to explaining number competences is a theoretical linguist, Heike Wiese (2007). There is of course a great deal more to mathematical competences than understanding and use of cardinal numbers, including topological and geometric competences that still (in 2015) have not been adequately modelled or explained.

Chapter 9 (partly inspired by the work of Max Clowes on human visual perception, whose ideas about "domains" I generalised by allowing more concurrently active interpretative domains, illustrated by the operation of the POPEYE program) explained in outline how visual perception could make use of a combination of bottom-up, top-down, and middle-out processing, straddling a variety of structural domains (not necessarily based on the kind of 3-D to 2-D projection proposed by Marr and others as essential to biological vision). Unlike some of the promoters and defenders of AI I thought that many of the problems were very difficult, especially problems in machine vision. Section 9.12 ended with the remark

"I do not believe that the progress of computer vision work by the end of this century will be adequate for the design of domestic robots, able to do household chores like washing dishes, changing nappies on babies, mopping up spilt milk, etc. So, for some time to come we shall be dependent on simpler, much more specialised machines."

Such tasks still (in 2019) remain well beyond the competences of robots, and I believe that the current approaches to intelligent robotics based on vast amounts of statistical learning will merely give the *impression* of closing the gap between natural and artificial vision systems, without *actually* doing so. In particular, it will not allow robots to replicate the ancient forms of spatial reasoning that led to the mathematical discoveries reported by Archimedes, Euclid, Zeno and many others, or even the spatial intelligence of crows and squirrels. (I suspect that will require new forms of computation, a topic discussed in some online papers, e.g. this incomplete, highly speculative paper: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/super-turing-geom.html>)

Chapter 10 began to speculate about how the sorts of mechanisms discussed in the book and others being developed in computer science and AI might account for some aspects of human consciousness, including explaining how some aspects of an individual's mind may be inaccessible to that individual while others are not. The chapter built on some of the architectural ideas in Chapter 6 and later chapters. Later work developed the ideas in the context of the Cogaff project mentioned in that chapter. Further features of consciousness, including evidence for the existence of qualia were later explained in terms of layers of *virtual machinery* combined with meta-cognitive mechanisms, e.g. in Sloman and Chrisley (2003).

The Epilogue, a late addition anticipating by several decades some of the concerns about the so-called "AI Singularity", explained (semi-seriously) in outline how super-intelligent machines might improve our planet by limiting the freedom of humans, whose morals, knowledge and intelligence, on the whole, leave much to be desired. People who worry about "the singularity" (e.g. in 2019) don't understand how far current Artificial Intelligence lags behind human (and squirrel and toddler) intelligence in crucial ways. All the worries people have should be directed at the **humans** who design and deploy such machines, just as they should be concerned about all the other potentially dangerous machines (e.g. aeroplanes able to drop bombs) designed by humans.

It's above all the humans who design, deploy and use the machines that need to be controlled (or at least well educated), not the machines (at present).

The Postscript, another late addition, attempted to explain how a computer programming language, like human languages, could be a rich and powerful tool yet allow the specification of contradictions arising from useless non-terminating procedures analogous to the expression of Russell's Paradox (and others) in human languages, a point developed in more detail in Slovan (1971)

I have tried to indicate how most of the ideas linking AI and Philosophy (and directly or indirectly also psychology, neuroscience, and biology) discussed in the book are still under development, with many problems still unsolved. The Turing-inspired Meta-Morphogenesis project, summarised in the next section, vastly expanded the scope of the research.

The Meta-Morphogenesis Project (proposed in 2012)

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

My work since 1978 has addressed many examples of types of possibility that need to be explained by attempting to construct at least outline explanations or research strategies for seeking and testing explanations of possibilities. [Examples, with references, will be added here later, including possible functions of visual perception, possible varieties of motivational and emotional state, possible kinds of mathematical discovery, possible forms of evolutionary change, possible forms of development of individual minds, and many more.]

The most ambitious of version of this crystallized as a goal late in 2011, when I had been asked to contribute to a book being put together for the Turing Centenary, and the editors somehow decided that I should comment on Turing's 1952 paper "The chemical basis of morphogenesis", now one of his most influential papers among scientists. It is a fine example of a great scientist proposing an explanation of a class of possibilities (though I don't claim to have taken in all the mathematical details). But it led me to ask "What might Turing have done had he not died two years after publication of that paper?"

That question led me to propose [the Meta-Morphogenesis Project](#) as the answer to the question.

The project aims to identify the many changes in information processing produced by biological evolution and its products since the very simplest life forms (or pre-life forms), and to propose explanations of how they are possible, though at present that is beyond our reach for many of the important examples.

New developments in biological information processing do not occur in accordance with some predictive law specifying biological necessities or regularities. But the facts show that an enormous variety of possibilities has been realised, and if we understand how those forms of information processing came into being and what made them possible, that will help us see that what actually exists is part of a much larger realm of possibilities that science needs to explain. In some ways this is similar to, though in the long run far more complex than, the explanation of the possibility of a wide variety of chemical elements with systematically varying physical and chemical properties, provided by the work of Mendeleef, Moseley and others.

https://en.wikipedia.org/wiki/Henry_Moseley

https://en.wikipedia.org/wiki/Periodic_table_of_the_elements

One of the key ideas that has come out of that investigation is the idea of a construction kit: evolution and its products make use of many sorts of construction kit. And at any time the products of the construction kits that have so far developed can make possible the development of new kinds of construction kit. This is very obvious in the history of technology. But the depth, variety and complexity of the biological phenomena are far greater, and much less well understood -- especially the implication that natural selection is a process that makes and uses mathematical discoveries, albeit blindly, as discussed in this draft paper:

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

Biology, Mathematics, Philosophy, and Evolution of Information Processing

Closely related:

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

Key Aspects of Immanuel Kant's Philosophy of Mathematics Ignored by most psychologists and neuroscientists studying mathematical competences. (Also pdf).

Potentially related:

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

The Meta-Configured Genome: Multi-layered, multi-stage, epigenesis

Based on collaboration with [Jackie Chappell](#) (School of Biosciences)

Also potentially related

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

Biologically Evolved Forms of Compositionality: Structural relations and constraints vs Statistical correlations and probabilities

What is information?

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

Why Claude Shannon's notion of "information" is far less relevant to understanding minds than the very much older notion, used, for example, by Jane Austen in her novels, a century before Shannon.

Draft documents on the Meta-Morphogenesis project (a potentially huge project, still largely unnoticed), the role of construction kits, the implicit mathematical discoveries made by biological evolution and by its products, and links to further work in that area can be found in these documents and in online papers that they refer to:

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

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

THEMES FROM THE 1978 BOOK DEVELOPED IN SUBSEQUENT WORK

Themes to be added:

- Development of ideas about architectures, e.g. in the CogAff project at Birmingham:
<http://www.cs.bham.ac.uk/research/projects/cogaff/#overview>
- Development of ideas about the space of possible sets of requirements ("niche space") the space of possible designs ("design space") the mappings between the spaces, and evolutionary, developmental, social, and other trajectories through those spaces. Here's a sample of papers, showing how the ideas and modes of presentation developed.
http://www.cs.bham.ac.uk/research/projects/cogaff/Aaron.Sloman_explorations.pdf
Explorations in Design Space, *Proc. 11th European Conference on AI*, Amsterdam, August 1994

<http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#41>

Exploring design space and niche space, Invited keynote for 5th Scandinavian Conference on AI, Trondheim, May 1995

<http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#37>

The “Semantics” of Evolution: Trajectories and Trade-offs in Design Space and Niche Space. Invited talk at 6th Iberoamerican Conference on AI (IBERAMIA). 1998

<http://www.cs.bham.ac.uk/research/projects/cogaff/00-02.html#57>

Architecture-Based Conceptions of Mind (Final version) (Invited talk) in *Proceedings 11th International Congress of Logic, Methodology and Philosophy of Science*, Krakow, Poland,

- Ideas about affective mental states and processes: motives, attitudes, emotions, moods, preferences, likes, dislikes, obsessions, ambitions,
- Ideas about virtual machinery, supervenience, anti-reductionism, causation in virtual machines, how virtual machines and architectural features can explain what qualia are, etc.
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>
Virtual Machine Functionalism (VMF) (The only form of functionalism worth taking seriously in Philosophy of Mind and theories of Consciousness)
- The mythical turing test: why the notion of a tests for intelligence is absurd and was not proposed by Turing, contrary to popular opinion. What is really needed is a set of tests (evaluation criteria) for a **theory** of (varieties of) intelligence.
- What information is and is not, and why representations are required. Why simple definitions of "information" are misguided, and need to be replaced not by a better definition, but by a theory of information. The universe contains matter, energy and information.
Jane Austen's theory of information in *Pride and Prejudice*:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/austen-info.html>
- Work on computing in education and why teaching computing has a deeper value than merely providing industry fodder, giving kids fun, recruiting more students for computer science degrees, etc. E.g.
<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk105>
- Development of software tools and teaching materials, at Sussex University then Birmingham.
Poplog and the Pop-11 language
<http://www.cs.bham.ac.uk/research/projects/poplog/freepoplog.html>

The SimAgent toolkit

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

- **Mathematics**

Gaps in philosophy/foundations of mathematics and current AI models of mathematical competences. The need to understand the roles of mathematics in various aspects of evolution (long before there were human mathematicians). Evolution as a [Blind Theorem Prover](#).

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

Biology, Mathematics, Philosophy, and Evolution of Information Processing

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

Discussion of a method for trisecting an arbitrary angle in a plane surface, apparently contradicting the proofs of impossibility. This has implications for philosophy/foundations of mathematics.

Compare Chapters 7 and 8 of CRP

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap7>

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap8>

Since 1978, I have given a number talks and conference presentations on the nature of mathematics, and the roles of mathematical discoveries in human and animal intelligence. Several of the talks in this directory discuss aspects of that topic:

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

An incomplete, still growing, discussion of the role of "Toddler theorems" in child development is here:

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

- More references to be added.
-

TO BE REVISED AND CONTINUED

References

- <http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#61>
A. Sloman, (1984), Experiencing Computation: A Tribute to Max Clowes, in *New horizons in educational computing*, Ed. Masoud Yazdani, Ellis Horwood Series In Artificial Intelligence, pp. 207--219, Chichester.
(Now also includes an annotated bibliography.)
- Anthony J Leggett, 1987 *The Problems of Physics*, OUP, Oxford,
- John McCarthy and Patrick J. Hayes, (1969), Some philosophical problems from the standpoint of AI, *Machine Intelligence 4*, Eds. B. Meltzer and D. Michie, Edinburgh, Scotland, pp. 463--502, Edinburgh University Press,
<http://www-formal.stanford.edu/jmc/mcchay69/mcchay69.html>
- <http://dx.doi.org/10.1111/j.1746-8361.1978.tb01321.x>
Karl Popper, (1978) Natural Selection and the Emergence of Mind, in *Dialectica*, 32, pp. 339--355, Blackwell Publishing Ltd,
- T. Shallice and M. E. Evans, (1978), The involvement of the frontal lobes in cognitive estimation, *Cortex*, 14, pp. 294--303,
<http://www-personal.umich.edu/~evansem/shallice-evans.doc>
- A. Sloman and R.L. Chrisley, (2003), Virtual machines and consciousness, *Journal of Consciousness Studies*, 10, 4-5, pp. 113--172,
<http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200302>

- Heike Wiese, The Co-Evolution of Number Concepts and Counting Words *Lingua*, 117, 2007, pp. 758--772,
http://www.uni-potsdam.de/u/germanistik/fachgebiete/geg-spr/pdfs/Wiese2007_NumberCounting.pdf
-

CRP Online Book contents page

Created as a template: 29 Jul 2015

Updated: 10 Aug 2015; 1 Sep 2015; 23 Dec 2015
