



Family Resemblance vs. Polymorphism

A comparison: Wittgenstein's Family Resemblance Theory vs. Ryle's Polymorphism and Polymorphism in Computer Science/Mathematics

And perhaps Kant's notion of "schema"?

(DRAFT: Liable to change)

Aaron Sloman

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

Installed: 30 Apr 2011

Updated: 11 Jan 2019 (major additions re parametric polymorphism).

24 Mar 2017; 2018; 14 Jul 2018

30 Mar 2016; 16 Jul 2016; 12 Sep 2016; 18 Dec 2016

21 Mar 2016 added "Polymorphism of design requirements" and a few edits.;

9 Mar 2016 added "creativity"

19 Apr 2014; 24 Apr 2014; 23 Oct 2015 (Reformatted/minor additions);

This paper is

HTML:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html>

PDF:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.pdf>

A partial index of discussion notes is in

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

This is one of several papers related to the "Computational Qualia" project summarised here by Ron Chrisley:

<https://www.researchgate.net/project/Computational-qualia>

With thanks to Luc Beaudoin for a quarter century of useful comments and discussion

<https://cogzest.com/>

CONTENTS

- [CONTENTS](#)
 - [SUMMARY](#)
 - ["Measure boundary" concepts for domains with a metric](#)
 - [Use of task-related boundary](#)
 - [Metrical and non-metrical concepts, and polymorphism](#)
 - [Parametric polymorphism in biological designs](#)
 - [Family Resemblance](#)
 - [Ryle's Polymorphism](#)
 - [Kant and polymorphism](#)
 - [Parametric polymorphism in computer science](#)
 - [Exercises for the reader](#)
 - [Evolution and parametric polymorphism](#)
 - [Programming languages with parametric polymorphism](#)
 - [More examples of parametric polymorphism](#)
 - [CONJECTURE: Other concepts exhibiting parametric polymorphism](#)
(Added 25 Apr 2014)
 - [Placeholders for concepts to be analysed in terms of parametric polymorphism.](#)
(Added 27 May 2014)
 - ["Conscious" as an example of parametric polymorphism](#)
(Added 23 Oct 2015)
 - [Consciousness and Euclidean geometry](#)
 - [Polymorphism of design requirements](#)
(Added 21 Mar 2016: To be expanded)
 - [PHILOSOPHICAL LAMENT](#)
 - [REFERENCES](#)
-

SUMMARY

NOTE:

After converting from maths to philosophy (roughly 1959-1962) I encountered a game discussed by philosophers. Choose a word (e.g. "consciousness") and look it up in a good dictionary. Then look up the main words used in the definition. See how long you can continue before the dictionary takes you back to a previously encountered word. Prizes can be awarded for the smallest circle, or the largest circle without intermediate repetitions, or the longest trail to the start of a circle, etc. Naturally, the game was invented by a mathematician as explained in [https://en.wikipedia.org/wiki/Vish_\(game\)](https://en.wikipedia.org/wiki/Vish_(game)). I have the impression the game is often played unwittingly by philosophers.

Philosophers, linguists, and no doubt many others, have been puzzled by the fact that some words that we all use effortlessly in everyday conversation are very difficult to define without generating definitional circles. Examples include "goodness", "truth", "knowledge", "probable", "science", "life", "consciousness", "experience", "belief", "game", "rationality", "normality", "intelligence", "pile", "heap", "big", and many more. Some of them are harder than others, and the causes of difficulty are not all the same. "Conscious", "conscious", "experience", and related words are particularly

difficult.

This discussion paper briefly summarises some of the problems, and some of the solutions that have been presented by philosophers, and then attempts to relate the problems to deep ideas developed in Computer Science since the middle of the 20th century, building on older ideas in mathematics -- ideas that appear to have been ignored by most philosophers, especially the concept of "parametric polymorphism", discussed below. (I think it can be argued that the same concept is used, with different labels, by mathematicians, but will not pursue that now.)

I shall also say a little at the end about how the notion of "virtual machine functionalism", which is related to complex and largely unrecognized developments in computer systems engineering, is relevant to some aspects of what we call "consciousness".

The problems are not really about words considered as sequences of characters or sounds, but about *concepts*. I'll treat the words as corresponding to or expressing concepts that can be expressed by different words in different languages, and by synonyms in the same language. So this is a discussion not just about *words* but also the *concepts* they express (a notion that I shall not attempt to analyse here -- though it exemplifies the problems discussed here).

I shall also say nothing about the claim that no concepts (or words) can really be defined without implicit circularity or an endless regress since all definitions must use words that also need to be defined. I'll also ignore [Austin \(1939\)](#). Compare G.E.Moore's "[paradox of analysis](#)".

"Measure boundary" concepts for domains with a metric

Some of the problematic words clearly refer to divisions along some measurable dimension, e.g. "big", as applied to trees, rocks, or waterfalls, refers to *size*. But the problem is to specify what sizes make something big or not big. One popular theory is that there is always implicitly or explicitly some class of comparable objects that provide the standard, as in "big house", "big London house", "big animal", "big elephant", "big flea", etc. But that does not explain where the boundary between big and non-big members of the class should be. A common answer is to take the average size, or the most common size, or some other statistically defined size as the boundary.

(Note: Some psychologists seem to think we learn where the boundaries should be by being trained by more experienced language users. [REFS] That begs the question of why those boundaries were originally chosen by intelligent humans!)

I suspect, but will not argue here, that there are normally principled explanations for choices of (possibly fuzzy) boundaries that have nothing to do with statistics, but are more concerned with phase transitions in functional relationships (e.g. affordances changing from positive to neutral to negative), as illustrated in

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

If there were a common rule, something like "X is a big member of class A if and only X is bigger than the average size of members of A", where size is measured in some well defined way for members of class A (e.g. using height for humans area for lakes, and weight or passenger capacity for airliners), then the word 'big' would no longer be vague or indeterminate and the sentence "X is a big A" would be an example of parametric polymorphism.

If the selection of dimension for size uses no general principle and speakers of the language merely have to learn which dimension is used for each class, then the polymorphism is partly *ad hoc*.

But there may be a principled method for determining the dimension. In some cases of natural objects it may be the perceptually most salient dimension e.g. physical volume in the case of elephants. In connection with objects designed by humans there may be one or more physical measures that designers are most concerned with maximising (or minimising), e.g. speed, reliability, comfort, etc. In such cases, including human designed transporters (cars, boats, bicycles, airliners, etc.), the parametric polymorphism is *principled*. However, in some cases the dimension of variation referred to is partly *ad hoc* (e.g. sleekness in a car?) and partly principled.

(I suspect some empirical linguist has already studied ways in which many creative uses of language can depend on discovery of a new parametrisation of a familiar concept.)

In more complex cases the comparison with an average might be restricted to the average for some explicitly or implicitly specified subset of the class A: "X is a big French lake" might use comparison of area, but restrict the computation of average area to French lakes. In some contexts the comparison class which determines the size measure and the restriction, and possibly some modified version of the size measure, might be inferred from the context in a well defined way.

Note that the object X need not be a member of the class in question as indicated by use of a subjunctive form something like "would be", as in "X would be a big French lake", said of a lake that is not in France, but in a country with much bigger lakes than any in France.

Use of task-related boundary

Another way in which the intended meaning of "X is a big A" could be restricted is by a need to meet some practical requirement. If someone is trying to reach a high shelf and asks for a big stool the criterion of adequacy will not be a statistical comparison with a class of stools but whether a stool is high enough to enable the asker to reach the shelf by standing on it. Intelligent listeners will not need to have that explained to them when asked about the availability of a big stool. Similar practical requirements can eliminate the vagueness of words like "pile" or "heap" as explained below, if they too are treated implicitly as examples of parametric polymorphism. (Those who are familiar with the work of philosopher Paul Grice will recognize this as a form of what he called "implicature".)

There are well known paradoxes concerned with these concepts, including the "sorites" paradox, i.e.: one grain of sand does not make a heap, and adding one more grain of sand to a collection of grains that do not make a heap will not produce a heap. From that it seems to follow that no amount of sand could form a heap.

There are many concepts that share this sort of property, including relational concepts such as "near to", and more abstract properties such as being "efficient", e.g. at lawn-mowing, or being an "economical" report (i.e. using relatively few words) or an "economical proof", in mathematics. There has been much philosophical discussion about such concepts though as far as I know no consensus regarding what they do or how they do it. See for example:

<http://plato.stanford.edu/entries/sorites-paradox/>
http://en.wikipedia.org/wiki/Sorites_paradox

Many such concepts are used in contexts where there is some practical need, and what does or does not meet that need determines where the boundary is taken to be, by the people involved in some discussion or choice, or planning activity.

For example, consider the problem of holding down a tarpaulin on a windy day, using a heap of sand on each corner: there will be some sizes of heap that are clearly insufficient for that purpose and others that are clearly more than sufficient (at least for the usual strengths of gusts of wind in that location). In between the insufficient and excessive sizes there will be a region of gradual change in which the resistance against being blown away gradually increases. At some point a person or group of people involved in using sand may decide that the current weight of sand on each corner will suffice.

In such a situation there is no need for the language or the practical context to determine an exact boundary between what does or does not suffice. An inexperienced picnic-goer may claim to have added four heaps of sand to the tarpaulin that are clearly adequate for the current wind strength. Yet someone with experience of wind fluctuations at that location may judge that more sand is required at each corner.

Asking "At what level does it become true that there is a heap of sand on each corner" pre-supposes that every statement must have a corresponding *sharp* true-false boundary, whereas many of the concepts used in linguistic communication, in thought, and in practical reasoning work in ways that do not satisfy that condition.

This is part of the motivation for "Fuzzy logic", which uses the notion of "degrees of truth". I am not advocating the use of any metrical concept, merely arguing that in some contexts the true/false boundary is a matter for practical decision, with no sharp distinction between right and wrong answers. Sometimes the variation is along a metrical dimension. In other cases it may be between structured entities: e.g. designs for a bridge that needs to meet various criteria.

The demand for an answer to 'Is the statement true?' has less relevance to such uses of language than requests for answers to 'Have you finished?' 'Could some gusts prove too much for that weight of sand?' 'Would it be a waste of time to search for more sand to add?' Will that sort of truss resist corrosion by weather?

Often the answers to such questions will be matters of practical judgement. For further discussion of such cases see:

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

Spatial prepositions as higher order functions

In some contexts the problem of deciding whether a statement is true or not does not hinge on similar practical decisions but on which among a collection of alternative sets of criteria is being used. If one person claims that Car A is a better buy than Car B and another person regards Car B as better, then this may not be at all like disagreements over whether enough sand has been added to meet an agreed need. Instead, the requirements against which cars are being judged can differ in complex ways, e.g. if one person wants a car for use in racing competitions while the other wants a family carrier. Such differences produce more examples of parametric polymorphism, often involved in the use of "better" as explained below. It is also possible for both people to want a car to be used for both purposes, but with different emphases on the relative importance of the two purposes.

Metrical and non-metrical concepts, and polymorphism

There are other words or concepts for which it is hard to specify necessary, or sufficient, let alone necessary *and* sufficient conditions, of applicability, and which cannot be analysed in terms of a transition point along a dimension, for there need not be a measurable quantity involved as in "big", "tall", "heavy", etc. Two of the best known examples, discussed at length by Wittgenstein in *Philosophical Investigations (PI)* (sections 65-71), are "game" and "language", discussed further below. Other words, or concepts, that have proved problematic refer to abstract properties of people, or their thought processes, or products of their thoughts, including words like "emotion", "feeling", "want", "believe", "know", "understand", and "explanation". Gilbert Ryle described some of these concepts as "polymorphous", as discussed [below](#).

Evaluative concepts

Another class of words concerning which it is hard to agree on conditions for their applicability is the class of evaluative terms, including "good", "bad", "ought", "beautiful", "pleasant", "awe-inspiring", and many others. Some of these are often discussed in connection with moral philosophy, but they can also be used in other contexts, including aesthetic judgements, types of holiday resort, football teams or players, and many more. I think the meanings of most of these words depend on the more basic concept "better", partly analysed, with its connections to "good", "ought" and other words in Sloman ([1969](#)), ([1970](#)).

Reading reports on consumer products and services in consumer magazines or web sites, such as the UK Consumer Association magazine *Which?*, reveals that there are many contexts in which "X is better than Y" makes a factual claim that can be explained in terms of a purely logical (topic neutral) notion of "Better" that is applicable not just to two items, as in "this is better than that", but requires additional factors in the comparison to be made explicit, such as: (i) a reference condition C, (ii) a respect of comparison R, (iii) the circumstances of comparison (or situation of comparison) S, (iv) a subclass of comparison Z. Then if, P and Q are properties, "Better(P, Q, C, R, S, Z)" can be read as "Being P is better in respect R than being Q with reference to satisfying condition C, in circumstances S, for members of the class Z". For more details on this proposal see Sloman ([1969](#)).

This pattern seems to fit many of the conclusions in consumer reports, where being P and being Q refer to makes or types of item that consumers might wish to purchase, e.g. tools, foods, holidays, cars, etc. The paper proposed a generic method for deciding whether such a statement is true or false independently of what the particular items P, Q, C, R, S and Z are, and to that extent "Better" was defined as a *logical constant*. (The detailed analysis had some gaps and imperfections, that are not relevant in this context.) That analysis presents "Better" as a concept exhibiting parametric polymorphism, and then showed how other uses of Better, including non-factual uses, could emerge from that use.

What I am suggesting here is that there are many more concepts of ordinary language that have similar more or less complex types of parametric polymorphism, including "X is conscious of Y" and "X is conscious that Y". The paper on "Better" was published in 1969 long before I learnt about the uses of parametric polymorphism in programming languages and software design.

[Jump to CONTENTS](#)

Parametric polymorphism in biological designs

It seems that biological evolution (unwittingly) "discovered" the power of parametric polymorphism long before human software designers, architects and other engineers did. That's because important abstractions were "discovered" by evolution and then used by instantiating them in different ways to meet a variety of specific needs, e.g. in different species, or in the same individual at different stages of development, or in different mechanisms within the same organism. Some aspects of these processes are discussed in connection with biological evolution's use of [compositionality](#), in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sloman-compositionality.html> (Also [pdf.](#))

For more on this see the Meta-Morphogenesis project:

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

I'll show below how the idea of "Parametric Polymorphism" developed in Computer science provides an illuminating way of explaining how some of these puzzling concepts work, though we need to extend the concept now used in Computer Science. I think all of these ideas can be treated as special cases of "Context sensitivity" of meanings, based on (sometimes implicit) parameters, an idea that can be construed as an extension of parametric polymorphism as currently studied and taught, e.g. in linguistics and computer science.

Family Resemblance

In his *Philosophical Investigations (PI)* (sections 65-71) Wittgenstein famously attempted to demolish the naive theory that all words referring to some class of phenomena do so in virtue of a set of features common to all members of the class. He used the example of the class of "games" and concluded that instead of something common to all games there is only a network of partial similarity relationships linking games to other games so that all games are connected to others in the network, but not directly: "... we see a complicated network of similarities overlapping and criss-crossing: sometimes overall similarities ...". This is often referred to as the "Family Resemblance" theory, described in more detail in http://en.wikipedia.org/wiki/Family_resemblance and in <http://plato.stanford.edu/entries/wittgenstein/#Lan>

It would be a useful exercise to take his examples (and many more) and show how they can be understood in terms of parametric polymorphism, where the relevant abstractions and parameters are not made explicit in our language, but implicitly understood by users. (Like many features of human languages.)

[Digression]

I have always thought that this analysis merely showed Wittgenstein's inability to think about explanations at an appropriate level of abstraction, although he was outstandingly good at producing *examples* to challenge over-simple philosophical theories produced by others. When he proposed his own alternative theories he was not so good. Partial recognition of his limitations may explain why he tried to claim that the job of philosophy is not to solve problems by producing new theories, but to dissolve them by removing the temptation to treat them as hard or interesting problems. This led him to compare doing philosophy with "showing the fly the way out of the fly-bottle" (a kind of insect trap), in PI 309, ignoring the great advances of earlier philosophers

(including Hume, Kant, Frege, and others). But my primary aim here is neither to expound nor to criticize Wittgenstein. **[End Digression]**

Related notions:

Notions that are closely related to the idea of "family resemblance concepts" include "cluster concepts" (Gasking?) and concepts with "open texture"(Waismann).

Ryle's Polymorphism

Gilbert Ryle also recognized the difficulty in saying what is common to instances of a complex concept, e.g. "thinking", "practising", "being angry", "attending", "desiring", and others, and suggested that many such concepts are "polymorphous". This was a rather vague notion, possibly suggested by the fact that some physical substances can take different visible forms despite having a common underlying structure (e.g. carbon as graphite and as diamond). But as far as I know he never made that notion very precise. Ryle's ideas were endorsed by a number of other philosophers including, for example, [White \(1967\)](#) who firmly endorsed and elaborated the theory that many concepts, especially a variety of concepts referring to mental states, processes or events are polymorphous.

Note: 24 Mar 2017

An earlier version of this document suggested that Ryle's ideas about polymorphism were presented in [Ryle \(1949\)](#), but I have since learnt that I was mistaken, insofar as I don't think he used that label in the *The Concept of Mind* even though he implicitly used the concept. It's possible that Ryle knew about and was influenced by the role of polymorphic type theory in *Principia Mathematica* (Russell and Whitehead). Anyhow the importance of polymorphous concepts for philosophy of mind and philosophy of language is independent of what Ryle thought or wrote.

Note added 15 Apr 2016

I have recently learnt (or perhaps been reminded) that in [Austin \(1959\)](#) Austin had referred to "adjuster" words, apparently expressing the same idea as Ryle and White expressed using the label "Polymorphous". Namely by means of adjuster words "other words are adjusted to meet the innumerable and unforeseeable demands of the world upon language". So what is meant by "real" (and we could also say the same about related words like, "genuine", "illusory", "fake", etc.) is that what is being said by the use of the word depends on another word or phrase with which it is being combined, as in "real pig", "real policeman", "real proof", "real conviction", etc.

A real pig and a real mathematical proof need have nothing in common, whereas a real pig and a fake pig, will have quite a lot in common (e.g. visible appearance from a distance) and real proof and an invalid proof may also have a lot in common, if the only error in the invalid proof is a something like a subtle omission of a case, somewhere in the middle.

The general idea that I have here labelled "polymorphism", following Ryle, has probably been rediscovered many times and given different names.

NOTE: The adjectives "polymorphous" and "polymorphic" are used interchangeably in English.

I think Ryle's partly understood notion of a polymorphous concept was very important. Something like that idea was later re-discovered and made far more precise in the second half of the twentieth century by mathematical logicians, programmers and computer scientists. In doing so they built, wittingly or unwittingly, on Frege's work in the previous century on generalising the concept of a mathematical function to include higher order functions, and work by [Church](#), [Kleene](#) and others, developing higher order logics and [the Lambda Calculus](#), and also [Typed versions of Lambda Calculus](#).

Kant and polymorphism

I suspect that Kant had something like this notion of polymorphism in mind when he criticised the common claim (e.g. made by Bishop Berkeley) that understanding a concept was something like perceiving and then remembering an image, or composite image, that could be used later to recognize instances of the concept. Kant opposed this with the idea that instead of attaching names to objects or images, intelligent individuals needed to learn rules for discriminating or identifying instances of the type named. Using rules to cope with generality included a rule for identifying triangles by specifying the number of corners and the number of sides while leaving the absolute or relative sizes unspecified. (He sometimes used "schema" instead of "rule".)

Because Kant was writing long before the development of programming technology he thought that our ability to use such schemata or rules, was very hard to explain, though in fact he almost invented the notion of a parametrized program:

"In truth, it is not images of objects, but schemata, which lie at the foundation of our pure sensuous conceptions. No image could ever be adequate to our conception of a triangle in general. For the generalness(generality?) of the conception it never could attain to, as this includes under itself all triangles, whether right-angled, acute-angled, etc., whilst the image would always be limited to a single part of this sphere. The schema of the triangle can exist nowhere else than in thought, and it indicates a rule of the synthesis of the imagination in regard to pure figures in space." (This idea is linked below, to parametrised concepts.)

"The conception of a dog indicates a rule, according to which my imagination can delineate the figure of a four-footed animal in general, without being limited to any particular individual form which experience presents to me, or indeed to any possible image that I can represent to myself in concreto. This schematism of our understanding in regard to phenomena and their mere form, is an art, hidden in the depths of the human soul, whose true modes of action we shall only with difficulty discover and unveil."

Immanuel Kant *Critique of Pure Reason*, 1781.

http://www.philosophy-index.com/kant/critique_pure_reason/i_sf_ii_i.php

How would Kant have responded to the suggestion that the rule or schema accepts parameters in such a way that the same rule combined with different parameters can be used to generate, or classify items in different categories where what is common to members of a category may be a set of parameters? I suspect he would have liked the idea, as enriching his concept of a "schema". In fact his geometrical examples, e.g. "triangle" "quadrilateral", etc. are obvious uses of a schema with parameters. Showing how "polygon" fits this idea is left as an exercise for readers!

Is the generalisation of "triangle", "quadrilateral", etc. to "polygon" comparable to generalisation of "cat", "dog", "horse", "mouse", to "species"? Could Kant have grasped something like this without having the vocabulary to express it? (Hence his reference to "an art hidden in the depths of the human soul", perhaps?)

In more complex cases there may be hierarchies of such categories, including, after additional abstraction, the concept of a *species*. This idea is essentially Frege's discovery that predicates (including relations) can be thought of as functions that can take parameters and produce values, where the parameters and values do not need to be numbers, as in the original mathematical functions. In particular, higher order functions can take functions as parameters, and that is also true of many higher order concepts, as illustrated in the discussion of "better", and other evaluative words, [above](#).

We now know that all this invention of new terminology and use of it in an extension of our familiar language is not enough: there remains the question of how such things can actually work? What mechanisms do brains need in order to enable them to perform the operations I have described (albeit vaguely). I doubt that any philosopher, psychologist, brain scientist or linguist could have expressed this precisely in the first half of the twentieth century. But since then advances in computer science, and disciplines influenced by computer science have made these ideas familiar and demonstrably useful insofar as they extend the powers of programming languages (and their interpreters, compilers and run-time support systems) to enrich the space of designs for working computer systems, including new models of working minds. (Not yet with the complexity of human or squirrel minds, however.)

And Kant was dead right in saying that merely adding new kinds of abstract images to the run-time systems could achieve the same result. Kant came close to inventing key ideas of Artificial Intelligence.

Parametric polymorphism in computer science

Leaving aside what Ryle may or may not have intended, I think the key idea that is missed by the family resemblance theory, but was partly understood by Kant, is that things that do not have any shared properties (such as height, colour, number of components, habitat, visible behaviour) may nevertheless have something more abstract in common which can be described in terms of a **common schematic description**, that contains gaps (in the sense in which function symbols contain gaps according to Frege), such that when the common description is suitably instantiated (by plugging the gaps in various ways) it can produce, different, possibly very different, instances of a generic concept *in a principled way*, which in many cases can be expressed as a generic computer program for creating or recognising instances of the abstraction, or combining it with other abstractions to create new useful abstractions.

For example, consider what might be meant by a biologist describing an animal as an "efficient forager". This label could be applied to animals of very different types (living in water, on land, in mud, in air, in the bark of some plant, etc.) with very different physical forms and very different physical movements.

What they have in common that justifies calling them efficient foragers might be that they form instantiations of O in something like the following abstract description:

O is a living organism and O consumes physical matter of type F (Food), which O obtains by performing actions A which usually bring O to instances of F and O needs to consume at least quantity Q of F per time period T, and the actions A have the effect that O regularly consumes at least Q of F in each period T.

In some cases this would have to be made more complex to allow for O requiring different types of food at different times, or when internal sensors indicate different deficiencies.

The different organisms that satisfy this description (when instantiating O) may require very different types of F, and perform very different actions A, and need very different quantities Q per period T. So the appearances, sizes, habitats, and visible behaviours of an earth-worm, an insect, a fish that feeds on algae and a grazing mammal may all be very different (especially perceptually different) and yet they all satisfy the above schematic description of "efficient forager" with all the variables other than O either existentially quantified or determined by context.

Concepts that have such complex relational conditions for their application include not just concepts of particular mental states (e.g. being conscious, having an intention, being in an emotional state) but also concepts used in engineering design, including concepts defining meta-functional requirements ([Sloman and Vernon\(2007\)](#)), and concepts related to the design of sophisticated forms of virtual machinery [Sloman \(1993\)](#).

I think Ludwig Wittgenstein's theory of "family resemblance concepts", Friedrich Waismann's concept of "open texture", and many theories concerning vagueness of meaning, all fail to take account of parametric polymorphism as a deeper explanation of what's really going on, with many well-defined working instances now available in complex computer programs.

Exercises for the reader

Thinking about these examples may help puzzled readers get a deeper understanding of what has been said above.

(1) explain how interpreting 'fierce predator' as a concept with parametric polymorphism, can allow 'O is a fierce predator' to be true of very different types of organisms with very different sizes, shapes, prey and behaviours;

(2) explain how very different control mechanisms controlling many different processes in very different organisms, in many different environments, can all be usefully described as using 'homeostasis', or control by means of a negative feedback loop, with either a fixed or a variable target state. Control theory is full of examples of parametric polymorphism even though that label is not used. I think Norbert Wiener fully understood the importance of this when he wrote his [1948 book](#).

See also the discussion of Polymorphism in section 2 of

A. Sloman, (2010), Phenomenal and Access Consciousness and the "Hard" Problem: A View from the Designer Stance, Int. J. Of Machine Consciousness 2, 1, pp. 117--169,

<http://www.cs.bham.ac.uk/research/projects/cogaff/09.html#906>

[Jump to CONTENTS](#)

Evolution and parametric polymorphism

The power of natural selection in biological evolution seems to depend in part on the fact that it uses forms of representation and ontologies that support the specification of designs with parametric polymorphism, so that re-usable design patterns can be employed with different parameters in different species or in different stages in the evolution of one species. The wide-spread occurrence of homeostatic mechanisms in biology may be an example.

Mechanisms for acquiring, storing, interpreting, using, combining, deriving, modifying, and exploring many different kinds of information may provide additional examples. For example, many forms of perception of structures and processes, along with goal formation and learning mechanisms, seem to be used by organisms dealing with environments containing complex structures and processes whose existence does not depend on their being sensed. In particular, I have argued that many non-human animals, and also pre-verbal children have competences that seem to require the use of internal information structures that allow acquisition, manipulation and use of information items that (a) vary in complexity, (b) vary in structure, and (c) make use of compositional semantics (so that information contents can be expressed by newly constructed information bearers in a systematic way). These systems share so much functionality with human languages, despite not necessarily being made of linear strings (and human sign languages are not either), that they can be called 'Generalised Languages' (GLs). See [presentation](#)

Concepts for describing states and processes in organisms with such mechanisms will need to exhibit parametric polymorphism if they are to express the deep commonalities between superficially different organisms and their behaviours. I suspect this is true of many concepts referring to mental states and processes, as Ryle seems to have realised.

This discussion implies a type of decomposition of the genome that may not have been considered by biologists. Future research will be required to spell out the details of this process and support it with examples. This is related to the view of evolution as a 'blind mathematician' introduced here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/bio-math-phil.html> and the discussion of the Meta-configured Genome, here: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-configured-genome.html>

Programming languages with parametric polymorphism

This sort of schematic specification is commonplace in what are misleadingly called "object-oriented" programming languages, that allow definitions of "multi-methods" (sometimes referred to as "generic functions") that are applicable to objects of very different types. The label "parametric polymorphism" is often applied to method specifications where newly defined procedures have parameters, and what exactly the method does, or how it does what it does, depends, in a systematic way, on the combination of parameters supplied. A partial and highly simplified introductory overview of features of such programming languages can be found here: <http://www.cs.bham.ac.uk/research/projects/poplog/teach/oop>

More examples of parametric polymorphism

Other examples exhibiting this kind of parametric polymorphism include "X is an efficient Y", where Y can be instantiated with concepts of a type of entity with some sort of function and X can be instantiated by different examples of that type, and working out whether substituting a particular pair for X and Y produces a true statement can involve digging into particular features of the

concept referred to as Y, including whether it specifies a function, a state to be achieved, a state or property to be maintained, a process to be produced or prevented, and various additional conditions derived from what is substituted for Y are satisfied.

A fairly simple example would be X (a person) is a national of Y (a country). The criteria for nationality will be different in different countries, but there are systematic ways of finding out what they are for a particular country and checking whether X satisfies them. The concept of nationality does not specify what sorts of criteria should be used though one requirement is that Y has some sort of legal system. If a country is discovered whose people know nothing about law-making, law-enforcement, national boundaries, international agreements, etc. then it could be the case that whether Judy is a national of that country has no answer. The people may or may not already have institutions that allow that situation to be changed. If not, there may or may not be some well defined route to setting up such institutions.

I suspect there are far more concepts used in ordinary human languages that have this kind of polymorphism than has generally been acknowledged and that in many cases the parameters that specify the interpretation are not supplied explicitly in a linguistic form, but have to be inferred from context. This can lead to mistaken theories about the vagueness or indeterminacy of the concepts, and also to what I regard as false theories of meaning such as Wittgenstein's "family resemblance" theory -- unless his formulation of that theory and the role of criss-crossing similarities, etc. was intended as a first draft, not very well articulated, specification of what later came to be called parametric polymorphism in computer science, where the requirement for giving computers complex new kinds of functionality drove the conscious development of a pattern of concept specification that had already been unconsciously developed in cultures encountering a need for new more sophisticated types of word meanings.

I think there have been philosophers who, without using this terminology, and without knowing anything about programming and computer science, criticised Wittgenstein's theory by showing how concepts like 'game' and 'language' could be specified using something like the techniques for specifying polymorphic notions in programming languages. (I think I have read a few philosophical papers making this sort of point using different terminology, though I cannot now remember where.)

CONJECTURE: Other concepts exhibiting parametric polymorphism (Added 25 Apr 2014)

I conjecture that there are very many concepts that have proved difficult for philosophers to analyse because they are complex cases of parametric polymorphism. Examples include: "proof", "ability", "power", "theory", "cause", "evidence", "explanation", "design", "understanding", "mechanism", "teach", "fear", "emotion", "love", "surprise", and probably many more, including especially many concepts referring to mental states, processes or capabilities.

I also conjecture that in most cases, perhaps all of them, focusing on the associated abstract noun and trying to define it will lead to failure. Instead we need to study the many different forms of complete sentence (expressing assertions that could be true, goals that could be realised, and questions that could be answered). Those sentences will, in many cases use not the noun form of the concept label, but corresponding adjectives, adverbs, and verbs, often tolerating ellipsis because some parameters are easily derived from the context.

Then unravelling the varieties of ellipsis, in order to identify the different explicit (non-elliptical) forms of expression in which the concept C can be related to a set of parameters $C(X_1, X_2, X_3, \dots)$ may help us find principled ways in which particular parameters determine requirements for truth (or intention fulfilment, or questions to be answered, etc.),

We'll also have to show how statements not mentioning those parameters either implicitly assume particular parameters derived from the context, or else existentially quantify over the parameters. All that will constitute a hypothesis about the type of parametric polymorphism involved in the concept C. This is generally not an easy task, and I suspect that in many cases carrying out this project will require the resources of a programming language that allows the formulations proposed to be tested and demonstrated working in running computer programs. In some cases (e.g. concepts referring to complex mental states or functions) it may be necessary to use advanced Artificial Intelligence toolkits to implement systems complex to model the phenomena being investigated.

Often such tests will produce failures or other surprises that indicate that the proposed analysis is incomplete or incorrect in some way. Perhaps something like this could realize Leibniz' dream of resolving complex disputes by saying "Let us calculate", but with the aid of more powerful tools than pen and paper.

Placeholders for concepts to be analysed in terms of parametric polymorphism.

(Added 27 May 2014)

- "Better", "good", "ought", "should", "must"
See [Sloman \(1969\)](#), [Sloman \(1979\)](#)
- Spatial prepositions, e.g. "near", "above", "left of", "behind", "in", "on", "at", ...
See [Sloman \(2006\)](#)
- Words associated with knowing, coming to know, not knowing, etc., e.g. "know", "learn", "discover", "notice", "surprise", "forget".
- **Added 9 Mar 2016**
"creative", "creativity", for example the differences between being a creative dancer, a creative cook, a creative mathematician, a creative engineer, a creative philosopher, a creative scientist, a creative novelist, playwright, or poet, the kinds of creativity found in biological evolution, and in various kinds of biological organism, ...
Some aspects of creativity in evolution are discussed in connection with the meta-morphogenesis project:
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-morphogenesis.html>
and an incomplete draft document
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/creativity.html>
- Add connection with John Barnden's ATT-META system (which he may not regard as relevant).
- others to be added

[Jump to CONTENTS](#)

"Conscious" as an example of parametric polymorphism

Many philosophers, psychologists, physiologists, anaesthetists, neuroscientists, biologists, and others have attempted to define "consciousness". The result is considerable diversity and, what's much worse, no sign of any agreement on criteria for evaluating such definitions.

As a result much of the discussion about what consciousness is, what it does, what functions it has (if any) how it evolved, what sorts of things can have it... etc., is in my view too unfocused and devoid of any lasting value. There are even contradictions in the things the same person may say at different times, some of them noted in Sloman and Chrisley (2003).

Reports of empirical results are more useful, though even those are often spoiled by use of question-begging modes of description, or over-generalisations from narrowly controlled experiments to a wide range of phenomena.

I suggest -- but will not attempt to defend this suggestion here -- that we can learn much if we (a) focus on the adjective "conscious" (and related adjectives) rather than on the noun (or nouns), and (b) consider a wide variety of sentence forms in which that adjective can occur, such as

- X is conscious of Y
- X is conscious
- X is no longer conscious
- X is not conscious of Y
- X was made conscious of Y by Z (e.g. a feature of Y, an event involving Y, or something said about Y or done to Y by someone or something else).
- X is conscious that Y
e.g. conscious that someone is watching him, that he is growing more popular, that the bridge is creaking, that the table top appears to have a parallelogram shape when viewed from certain locations, etc.

including a wide variety of examples of these forms, expanded into more complex forms of various sorts, e.g.

-- X began to be conscious of Y when Z

-- Doing A made X conscious that P

-- Being conscious of Y made X feel Z (e.g. afraid)

-- ... and many more. (To be expanded later).

We can investigate the wide variety of cases (some obvious, some problematic) to be found among living things including perhaps microbes, houseflies, woodlice, sunflowers, and a wide variety of animals, and animal collectives (troops, colonies, hives, swarms, families, nations, etc.). We can also, in parallel with that, investigate the many and varied requirements for machines of various sorts especially robots with various kinds of sensory or perceptual systems, various mechanisms and competences for acting or communicating, various kinds of autonomy, various kinds of learning ability, various kinds of potential for further development.

This sort of investigation is in part an application of the techniques of "Conceptual Analysis", explained in Chapter 4 of [Sloman\(1978\)](#):

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap4>

CHAPTER 4: WHAT IS CONCEPTUAL ANALYSIS?

This will help us appreciate the many different **requirements** (niche pressures) that influenced biological evolution towards different ways of being conscious of something and benefitting from that capability.

Then, adopting the designer stance, we can consider what sorts of mechanisms, including information-processing mechanisms, can provide various sorts of functionality, especially multi-level virtual machines with complex interacting control functions. It wasn't possible for philosophers to think about these until very recently (since about the mid 20th century). And even now many of them have not learnt how.

Such an investigation was begun (in a very small way) in Sloman ([1978](#)) chapters 6 and 10, including the observation that organisms that have many parts that can act more or less independently, will sometimes be in situations where there is a need to adjudicate between incompatible action tendencies of different parts or subsystems, or to meet a need to coordinate actions involving different parts.

For example, types of organism that (as a result of previous evolutionary and developmental changes) can have tendencies to move towards food, to protect their young, to build shelters, to run or fly or swim away from predators, may have different such tendencies triggered simultaneously. In very simple cases it may suffice for conflicting tendencies to be dealt with by automatic mechanisms that handle conflicts on the basis of some scalar measure of importance (or a weighted mixture of importance and urgency -- time left before it's too late to act), whereas in more complex cases the options available, the reasons for action, the possible consequences need to be considered and used by some more or less sophisticated, and possibly trainable, 'central' decision-making system.

[Chapter 10](#) of Sloman ([1978](#)) conjectured:

"So, perhaps the distinction between what we are and are not conscious of at a particular time, is concerned with the difference between information which is made available to, or used by, central administrative processes, and information which is not. There will be many processes which continue without any notice being taken of them by the central administrator, and at each moment there is an enormous amount of unused information present in stores of various kinds. There is no point cluttering up the central decision-making with all the details of the sub-processes: the task of relating all the information would be too unmanageable. So censorship of a sort is a prerequisite for normal functioning of such a system, rather than an oddity to be explained." (This principle is integral to the design of the POPEYE program described in Chapter 9 of [Sloman\(1978\)](#) available here:

<http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap9>

Later work, e.g. Sloman and Chrisley ([2003](#)), and [Sloman \(2010\)](#) stressed the importance of the role of virtual machinery in supporting rapid reorganisation of information processing functions and self-monitoring of some of the temporary contents of those processes, including intermediate 2-D information structures used in constructing 3-D interpretations of a stream of visual sensory data (the origins of qualia of various sorts). See also the discussion of Virtual Machine functionalism [below](#).

It is widely assumed that any such system should have the function of maximising expected utility for the individual but I believe evolution has found many other ways to support conflict resolution and decision making, including *architecture-based* as opposed to *reward-based* motivational mechanisms, as described in Sloman ([2009](#)). That theory extends the concept of reflexes from reflex physical actions to reflex motive generation.

Some of those reflexes may be genetically determined, others strongly influenced by learning processes interacting with genetic mechanisms -- but not because the individual learns what is and is not rewarding. Instead evolution has already "learnt" which motive generators are likely to produce more offspring if added to the pool! Motives generated by reflex generators can compete with motives derived from expected positive and negative rewards. This is an indication of some of the features of motivation that make "having a motive" a concept with parametric polymorphism.

After investigation of a wide range of cases and exploring a variety of designs for which instances are built and tested, including designs for systems that can inspect some of their own internal information processing (e.g. perceptual contents, motives, conflicts, planning processes, learning processes, etc.) we'll be able to develop new, precise, sophisticated, well tested, systems of concepts, including concepts with parametric polymorphism to cope with diversity of contexts and domains of application.

I suspect such a research process will have as deep an impact on our ideas about being conscious of various things and related ideas of believing, wanting, learning, planning, deliberating, deciding, hypothesising, explaining, teaching, ... as the influence had by development of evolutionary biology and its use in categorising and explaining observed biological evidence on previous concepts of species and forms of life. (A well worn example is learning that for the purposes of scientific understanding it's better to treat whales, dolphins and some other fish-like animals as mammals rather than fish.)

Of course, as with biology, and flat-earth theories, there may always remain some philosophers, biologists and psychologists who insist on sticking with their pre-scientific concepts of mind, consciousness, etc. and will continue going round the old debates, while the others constantly increase their knowledge and depth of understanding.

A sample of this approach applied to a variety of emotions, including long term grief, which most emotion theories cannot accommodate, can be found in this invited paper for the journal of Philosophy Psychiatry and Psychology Wright et al. ([1996](#)) dealing with a variety of affective states, including types of long term affective state that most theories of emotion ignore.

Insofar as the richer theories developed will have to cope with complex and changing internal structures, including processes monitoring and modifying such structures, it is very likely that any philosophical or other theory concerning the nature of such processes, the mechanisms involved, how they evolved, how they develop in individuals and what they do, will have to make use of what we have learnt about the role of *virtual machinery* and *virtual machine functionalism* in the last seven decades or so, developments that have so far touched very few philosophers. See also [VM-Functionalism \(2013\)](#).

We'll also need to broaden the range of examples discussed beyond such things as perceptual experiences of properties of things. A type of example that hardly gets discussed, although Immanuel Kant was aware of its importance, is mathematical experience of structures and

processes as illustrated in this discussion of the hidden depths of 'triangle qualia', which seem to have played an important role in the development of human understanding of Euclidean geometry before there were mathematics teachers and textbooks:

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

Consciousness and Euclidean geometry Added October 2015

Recently (October 2015) I have been trying to show how the mathematical forms of consciousness leading up to the powerful and deep mathematical results reported in [Euclid's Elements](#) over 2000 years, especially those concerned with geometry and topology (though I'll later show how arithmetic fits in too) might be products of biological evolution that built on more wide-spread biological capabilities concerned with perception of affordances, but not just the affordances that James Gibson wrote about: a wider variety including affordances concerned with possible and impossible changes in the environment that need not relate immediately to the needs or capabilities of the perceiver. A very important feature of those capabilities is that they have nothing to do with abilities to make inferences about what is probable or improbable, or to use statistical information. They are much deeper, and much harder to model, yet there is evidence that such mathematical competences are not restricted to mathematicians, or even adult humans, since some are detectable in pre-verbal toddlers.

That topic is being developed in some detail in these and other discussion papers:

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

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

There are major differences between being conscious of geometrical constraints in a class of deformable structures and being conscious of the colour of a visible patch. Any theory that deals with the latter but not the former cannot be an adequate theory of ways in which humans can be conscious. (Another example of polymorphism.)

There are many philosophers who have noticed and been puzzled by the essential privacy of some ways of being conscious, and some who have even denied the existence of private qualia. In Sloman and Chrisley (2003) it is shown how such privacy of content could be a feature of certain kinds of virtual machinery. For more about the philosophical importance of virtual machinery see this discussion of Virtual Machine Functionalism (VMF), contrasted with Atomic State Functionalism (ASF): [VM-Functionalism \(2013\)](#)

Polymorphism of design requirements (Added 21 Mar 2016: To be expanded)

The notion of something being well-designed clearly exhibits parametric polymorphism insofar as it involves being better-designed than it might have been and "better" is polymorphic, as explained [above](#).

But there are different sorts of design requirements, leading to different sorts of comparison as to whether something is or is not better than another. Some of the issues are discussed in connection with requirements for complex control systems [Sloman \(1993\)](#), and a partial analysis of meta-requirements for designs (e.g. for software systems) [Aaron Sloman and David Vernon\(2007\)](#).

Concepts with the sorts of parametric polymorphism referred to here cannot be learnt by systems seeking regularities in sensory or sensor-motor data. E.g. what's common to a well designed sports car, a well designed kitchen mixer, and a well designed virtual machine architecture cannot be found in input-output signals of users of information about such commonalities. It seems that evolution discovered the need for such concepts before we did. (I think all this implies that methods of "deep learning" cannot give rise to concepts with parametric polymorphism.)

[Jump to CONTENTS](#)

PHILOSOPHICAL LAMENT

The ideas summarised here and developed in the various papers and presentations referred to may be hard for most philosophers to digest if their philosophical education has been restricted to reading and commenting on the writings of other philosophers and perhaps a few psychologists, neuroscientists and physicists, without any first hand experience of learning to design, build, test, debug and explain working systems, especially working systems that include virtual machines with concurrently active interacting components and multiple modes of sensing and acting on their environments, illustrating virtual machine functionalism, and systems using parametric polymorphism with multiple inheritance (not explained here).

Unfortunately, most of the "experts" who *have* had that sort of experience, e.g. in designing, building and testing working robots, have not learnt enough philosophy to make their efforts relevant to deep philosophical debates. Perhaps that will change as a result of the revolution now in progress concerning the teaching of computing in schools -- as long as the computing education provided is not too narrow, and philosophically naive.

<http://computingatschool.org.uk/>

<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk105>

There is a recent paper by a philosopher recognizing the need to build bridges [Kondor\(2015\)](#).

TO BE CONTINUED

[Jump to CONTENTS](#)

REFERENCES

Some related papers and discussions

(There is a vast amount of relevant literature. This is not the place for a survey.)

- J.L. Austin, 1939, Are there a priori concepts? *Proceedings of the Aristotelian Society*, Repr in *Philosophical Papers*, OUP, 1961, pp 1-22
- J. L. Austin, 1959, *Sense and Sensibilia*, OUP, Oxford,
- Paco Calvo, What Is It Like to Be a Plant?, *Journal of Consciousness Studies*, 2017, vol 24, 9-10, pp. 205-227,
<http://www.ingentaconnect.com/content/imp/jcs/2017/00000024/F0020009/art00012>

- An online English version of Euclid's Elements is here:
<http://aleph0.clarku.edu/~djoyce/java/elements/elements.html>
- That was, arguably, the most important, and most influential, book ever written, ignoring highly influential books with mythical or false contents. Unfortunately, this seems to have dropped out of modern education with very sad results.
- Zsuzsanna Kondor, (2015), Theoretical Controversies - Terminological Biases: Consciousness Revisited,
Studies In Logic Grammar and Rhetoric 41 (54), pp. 143--160, DOI 10.1515/slgr-2015-0025,
<http://logika.uwb.edu.pl/studies/download.php?volid=54&artid=54-09&format=PDF>
- Gilbert Ryle, (1949) *The Concept of Mind* Hutchinson, London.
- <http://www.cs.bham.ac.uk/research/cogaff/papers.html#1969-02>
A. Sloman, How to derive "Better" from "is",
Am. Phil. Quarterly, 6, Jan, 1969, pp. 43--52
- <http://www.cs.bham.ac.uk/research/projects/cogaff/05.html#200508>
A. Sloman, "Ought" and "better",
Mind LXXIX, 315, July, 1970, pp. 385--394,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/>
A. Sloman
The Computer Revolution in Philosophy, Philosophy, Science and Models of Mind,
Harvester Press (and Humanities Press), Hassocks, Sussex, 1978,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#18>
Aaron Sloman, The mind as a control system,
In *Philosophy and the Cognitive Sciences*, Eds. C. Hookway and D. Peterson,
CUP 1993, pp. 69--110,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/96-99.html#2>
I.P. Wright, A. Sloman and L.P. Beaudoin,
Towards a Design-Based Analysis of Emotional Episodes, in
Philosophy Psychiatry and Psychology, 3, 2, 1996, pp. 101--126,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#200302>
A. Sloman and R.L. Chrisley, Virtual machines and consciousness,
Journal of Consciousness Studies, 10, 4-5, 2003, pp. 113--172,
- <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0605>
A. Sloman Spatial prepositions as higher order functions:
And implications of Grice's theory for evolution of language.
Online discussion note 2006 (with updates).
- Aaron Sloman and David Vernon, 2007, A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots, Contribution to euCognition wiki, (and CoSy Project technical report).
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/meta-requirements.html>

- <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/sorites.html>
A. Sloman
Sorites -- A New Analysis Based on Context
- <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/architecture-based-motivation.html>
A. Sloman, Architecture-Based Motivation vs Reward-Based Motivation,
Newsletter on Philosophy and Computers, American Philosophical Association,
Newark, DE, 09, 1, 2009, pp. 10--13,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk111>
Aaron Sloman,
"Evolution of minds and languages.
What evolved first and develops first in children:
Languages for communicating, or languages for thinking (Generalised Languages: GLs)?"
PDF Presentation 2008,
School of Computer Science, University of Birmingham,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/10.html#sab>
A. Sloman,
How Virtual Machinery Can Bridge the "Explanatory Gap", In
Natural and Artificial Systems, Proceedings SAB 2010, LNAI 6226,
Eds. S. Doncieux and et al., Springer, Heidelberg, 2010, pp. 13--24,
- <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>
Aaron Sloman
Virtual Machine Functionalism
(The only form of functionalism worth taking seriously in Philosophy of Mind)
Online discussion note (2013 + updates)
- Norbert Wiener,
Cybernetics: or Control and Communication in the Animal and the Machine,
The MIT Press, Cambridge, MA, (1948) 2nd edition, 1961,
- Alan R. White, 1967, *The Philosophy of Mind*, Random House, New York,
-

Maintained by

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