

Two Notions Contrasted: 'Logical Geography' and 'Logical Topography'

**Variations on a theme by Gilbert Ryle:
The logical topography of 'Logical Geography'.
Aaron Sloman**

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

This file is <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/logical-geography.html>
Also available as messy, automatically generated PDF file
<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/logical-geography.pdf>

Installed: 15 Oct 2006

Updates:

24 Feb 2016 (reformatted). March/August 2017 (Minor re-wording)
13 Apr 2011: corrected "Logical topology" to "Logical topography" (Thanks to Luc Beaudoin).
3 Jan 2010: added [link to work of David Jhave Johnston](#), and a response to one of his comments.
19 Aug 2009: fixed some typos and elaborated a few of the examples.
26 May 2009: added picture showing logical geography and topography
17 Jun 2008: added Portia spider example.

8 Jan 2008:

On re-reading Appendices III and IV of my Oxford DPhil thesis "Knowing and Understanding Relations between meaning and truth, meaning and necessary truth, meaning and synthetic necessary truth", (1962), online here: <http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1962>
I found to my surprise that several of the topics discussed below had been discussed in those two Appendices, and that the distinction between logical topography and logical geography was made in Appendix IV, discussing conceptual analysis, even though I did not use these words.

22 Oct 2007

(I had previously reported, from memory, Lakatos distinguishing 'progressive' and 'degenerative' research programmes. I was mistaken: he used 'degenerating' for the second category, though many authors present his ideas using 'degenerative'. I doubt that the distinction is of any significance, but I've altered my text to fit what he wrote.)

APOLOGETIC NOTE

On 4th January 2007 I announced to a philosophy mailing list that I was attempting to articulate a distinction between Gilbert Ryle's notion of 'logical geography' and a related but different notion of 'logical topography'. The letter is here:

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

Several philosophers kindly commented in ways that made me realise I had to make this a more complete and detailed discussion. In particular more than one of them suggested that I was simply making a distinction between empirical truths (logical topography) and conceptual (or analytic) truths (logical geography). Since I was already familiar with that distinction (I had written about

it years ago in this 1965 paper:

<http://www.cs.bham.ac.uk/research/projects/cogaff/62-80.html#1965-02>)

I was making a different distinction, I had to explain why that was not a good way to view what I was referring to. Unfortunately, elaboration of that point has made this document much longer and it still needs further revision.

NOTE

This is a draft discussion paper liable to change.

Comments and criticisms welcome. If you send me comments and do not wish to be quoted please say so.

If you make a copy of this please preserve a link to

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

to ensure that you access the latest version.

I'd like to thank Julia Tanney for help in understanding Ryle's ideas.

(Removed defunct link to her at University of Kent.)

Summary and request for comments

This paper distinguishes two versions of Ryle's notion of 'logical geography'.

- **Logical geography:**

The network of relationships between current uses of a collection of concepts.
(Probably what Ryle meant by the term.)

- **Logical topography**

Features of the portion of reality, or types of portions of reality, related to a given set of concepts, where the reality may be capable of being divided up in different ways using different networks of relationships between concepts.

Studying/analysing logical topography includes evaluating the alternative divisions in terms of the theories or applications that are supported by the different sets of concepts, or perhaps some other uses of the concepts.

NB: as the mathematical examples below show (e.g. notions of '[circle](#)' and '[ellipse](#)', and notions of '[number](#)'), the underlying reality in question need not be some physical reality. It can be anything about which we have a systematic theory that generates a variety of possibilities. (As discussed in Chapter 2 of *The Computer Revolution in Philosophy*: <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap2>)

(3 Jan 2010: Note stimulated by [comments by David Jhave Johnson](#))

The example of a functioning multi-component information-processing architecture (and many other examples from nature, e.g. weather systems, ecosystems, social systems, human minds) show that the underlying reality need not be static, as the example of the periodic table of the elements is.

Comments, criticisms and suggestions for improvement invited.

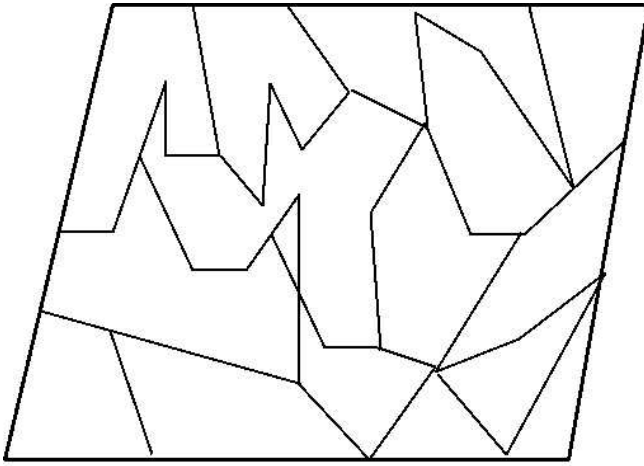
CONTENTS

- [Summary and request for comments](#)
- [Introduction: Ryle's notion of 'logical geography'](#)
- [Brain-Mind relationships](#)


- [Ryle's concept of 'Logical geography'](#)
 - [From actual to possible usages \(Added: 11 Feb 2007\)](#)
 - [Is study of logical geography empirical linguistics?](#)
 - [Conceptual analysis as deep science \(Updated 11 Feb 2007\)](#)
 - [A harder non-empirical task: Abduction of a new theory](#)
 - [Another kind of logical geography? Logical topography.](#)
 - [Do dynamical systems with attractors have motives?](#)
 - [Arguing about which definition is right, or best, is usually pointless](#)
 - [Problems of 'cluster concepts'](#)
 - [Problems of inconsistent concepts \(added 22 Oct 2007\)](#)
 - [Some ways of reaching agreement](#)
 - [Logical topography of Meta-Requirements for Robots](#)
 - [Logical topography of concepts of deliberative agents](#)
 - [Some earlier work on logical topography of 'better', 'good', 'ought', 'identity', etc.](#)
 - [NOTES](#)
-

Introduction: Ryle's notion of 'logical geography'

In his 1949 book [The Concept of Mind](#) (CM) and in other writings the philosopher [Gilbert Ryle](#) suggested that a good way for philosophers to resolve some philosophical disputes (often by discovering that both sides were based on conceptual confusions) is to study the 'logical geography' of the concepts involved. I used to think I knew what that meant. But now I think I was using a different concept of from Ryle's -- referring to a different type of analysis that is fundamentally related to the scientific project of providing explanations of how the world works. This paper provides some background then describes the difference I have in mind, showing how a theory about how some class of objects works generates a set of possible types of states, events and processes that can be referred to as a "Logical Topography". There are different ways of carving up that space into different categories or identifying different relationships that can occur within it. Those different ways define different "Logical geographies". This paper shows how work in AI and robotics can extend the work of Ryle and other philosophers by exposing logical topographies that support different possible logical geographies, only one of which may correspond to how our ordinary concepts work. This helps to resolve some century-old puzzles about the nature of conceptual analysis and to show how the relationships between philosophy and science can be deeper than many philosophers realise.



A Logical Geography
Map showing the structure of
a set of current concepts



A Logical Topography
Map showing the structure of
the corresponding part of
reality as revealed by scientific
and other empirical and
technological advances -- subject
to change as more is learnt.
One logical topography may
support several possible logical
geographies.

This paper provides some background then describes the difference between studying logical geographies and studying logical topographies, asking whether the distinction is a clear one and whether anyone knows if Ryle really did notice only one side of it.

Added: 30 Dec 2007

I believe this is related to Carnap's distinction between internal and external metaphysical questions.

I met Ryle when I was a graduate student gradually converting from mathematics to philosophy in Oxford (1957-1962) and attended some of his lectures and graduate seminars, though I was never formally his student. He did once kindly give me a 'tutorial' after I had submitted a paper to MIND, 'Rules of inference or suppressed premisses?'. He invited me to rewrite the paper 'in prose rather than the style of a railway guide' after which he published it in 1964. I've recently installed the published version [here](#). I now think it is a simple example of what this discussion paper is about, namely investigation of the 'logical topography' underlying a set of concepts, in contrast with what Ryle called the 'logical geography' of a set of concepts actually in use, which represents just one way of carving up the space of possibilities. However, I never had the chance to discuss his notion of 'logical geography' with him and don't recall having heard him talk about it.

At that time I misunderstood Ryle's theories about mind as being behaviouristic, and denying the existence of mental states and processes. That was partly because I mistakenly believed what other philosophers said about him, including philosophy students at Oxford who had picked up some of Ryle's turns of phrase but not what lay behind them.

Much later (in the 1970s) I realised, on carefully reading bits of CM (e.g. the chapter on 'Imagination') that by 1949 he had discovered the same notion of '[information processing virtual machine](#)' as began in later years to emerge in computer science, software engineering and AI, which was also crucial to [my 1978 book](#) (and much of the work of [Daniel Dennett](#)).

But Ryle did not have adequate conceptual tools to explain the ideas properly, e.g. because he wrongly thought (unless I've misunderstood him) that a mind-like virtual machine could be defined by a collection of externally observable behavioural dispositions of arbitrarily high order (i.e. dispositions to produce behaviours that are changes of dispositions toetc.), an idea that breaks down when there are reasons why there's not enough motor bandwidth for behaviours to be capable of providing information about everything going on inside, even over arbitrarily long times.

NOTE:

In many animals visual and other inputs add up to very much higher bandwidth totals than the bandwidth available for motor signals. So the full richness of the processing of sensory information can never be displayed in external behaviour. It's not even obvious that there is a closed, finite or infinite, set of true counterfactual conditionals regarding behaviour that could possibly fully characterise what's going on in a virtual machine made up of many concurrently and *asynchronously* interacting components, some of which may most of the time have no causal links with external behaviours though under some circumstances such connections exist e.g. if other components that do have access to effectors occasionally interrogate them.

For related reasons the contents of consciousness cannot possibly keep up with everything that's going on.

The links between information-processing in an internal virtual machine and behaviour can be very remote when the virtual machine is composed of several concurrently active, asynchronously interacting virtual machines. Moreover, in such cases, there is no neatly structured hierarchy of dispositions to produce external behaviour. Most of the interesting dispositions are dispositions to produce or modulate *internal* behaviours, including changing other dispositions. I think Ryle acknowledged such things, e.g. in his chapter on 'Imagination'. But I don't think he was able to think as a modern designer of computer operating systems would, about the variety of such dispositions and the number of internal states and processes that could be simultaneously active. (Even now, philosophers are not trained to think about such things, alas, so that many of their discussions about mind and body consider far too few possibilities.)

The logical topography (in the sense defined below) of that collection of ideas is what I've been working on for about 35 years, namely the analysis of the space of possible events, states, processes, and their components, in virtual machines with architectures rich enough to accommodate all known human (and some animal) mental phenomena. That space is still largely uncharted territory. Our actual concepts, whose logical geography carves up only a small subset of that space, is based on only very shallow and restricted understanding of the space. Nevertheless it seems to be a deeper understanding than any other type of animal has.

NOTE:

These points are related to the objection to a functionalist philosophy of mind based on the notion of a finite state machine, as discussed in this presentation on information processing virtual machines:

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

It is wrongly believed by many that if mechanisms do not affect behaviour while they are in operation they cannot be the results of evolutionary selection. This is wrong because a constantly active subsystem that only occasionally connects with external behaviour could have great biological advantages, even if most of its activity most of the time produces no biological advantage. A constantly active alarm sensing system that rarely gets used would be an example. For more on this see [The Mind as a Control System](#) (1993)

Brain-Mind relationships

Ryle was also not able to say much and wasn't interested in saying anything, about the relationships between entities, states, and processes in the virtual machine and those in the physical *implementation machine* -- the topic of neuroscience and neuropsychology. We still don't know how to explain the connections for most of the relationships between minds and brains, though we know a lot about other much simpler cases, e.g. the relationship between a running operating system (a virtual machine) and the underlying physical machine, or to be more precise, the underlying digital electronic (virtual) machine. (The underlying physical machines still hide mysteries physicists are trying to unravel.)

Anyhow, this paper is not directly concerned with his work on the Concept of Mind so much as with his notion of 'Logical Geography' which I am trying to contrast with a notion of 'Logical Topography', which I believe is a deeper notion.

Ryle's concept of 'Logical geography'

Ryle often used the phrase 'logical geography' to describe what he was doing, e.g. in the Introduction to CM:

'The philosophical arguments which constitute this book are intended not to increase what we know about minds, but to rectify the logical geography of the knowledge which we already possess.'

(Page 7)

(It is arguable that there is an inconsistency, since such rectification must in some sense 'increase what we know', but we can ignore that for now.)

For many years I have been referring to some of what I do (as a philosopher) as being clarification of what Ryle called 'logical geography', although until now I haven't used the term nearly as frequently as some of his admirers.

I have recently realised that there are two importantly different ideas, and I think Ryle was (mostly) using one of them, whereas I have been (mostly) using the other. It is possible that [Dan Dennett](#) (who was supervised by Ryle soon after I left Oxford) is probably somewhere in between. I discovered the mismatch as a result of reading an excellent introductory overview of Ryle's ideas by [Julia Tanney](#) available here: [Rethinking Ryle](#).

E.g. Tanney writes:

"....And Ryle's reminders of how these concepts are used, what we normally appeal to when we wish to defend or explain this use, and of what we need them for, are therefore as important now as they were fifty years ago. This investigation of their use - a cartographic exploration of the logical geography of expressions in which these mental concepts figure - may, in the end, tell against a proposed theory about the nature of their putative referents."

....

"Just as the recipe theorist must devise recipes by construing the role that ingredients, measures, and mixing and heating procedures play within a range of possible (successful) dishes, the philosopher's chart of the logical geography of concepts deals with the various ways in which these concepts figure in the sayings (not only the describings) of people competent in their use. Like the recipe-writer, the philosopher-cartographer will presuppose many abilities of the follower of his philosophical map."

....

"The chef in our thought-experiment knows how to produce delicious meals but is unable to abstract, generalise, and idealise what he does; for this, a recipe-writer is required. We may operate with expressions

down a limited set of tracks but will never know all the "implication threads" of the expressions we ordinarily and naively rely upon and operate with; for this a philosopher-cartographer is needed."

What I now think he meant by 'logical geography' was what Tanney describes, namely the network of relationships between a collection of connected concepts -- as they are *currently* used, where the network is described in a way that makes explicit actual connections and features of the usage of those concepts that may not be immediately evident to the users, and may even contradict what the users believe about their usage.

Ryle's paper "Ordinary Language", originally published in *Philosophical Review* LXII (1953) is now available online here <http://www.hist-analytic.org/Ryle.htm>.

(I think there are several of transcription errors in this version, possibly produced by an OCR package.)

The paper starts:

Philosophers' arguments have frequently turned on references to what we do and do not say or, more strongly, on what we can and cannot say. Such arguments are present in the writings of Plato and are common in those of Aristotle.

In recent years, some philosophers, having become feverishly exercised about the nature and methodology of their calling, have made much of arguments of this kind. Other philosophers have repudiated them. Their disputes on the merits of these arguments have not been edifying, since both sides have been apt to garble the question. I want to ungarble it

and ends

We are interested in the informal logic of the employment of expressions, the nature of the logical howlers that people do or might commit if they strung their words together in certain ways, or, more positively, in the logical force that expressions have as components of theories and as pivots of concrete arguments. That is why, in our discussions, we argue with expressions and about those expressions in one and the same breath. We are trying to register what we are exhibiting; to codify the very logical codes which we are then and there observing.

This can be compared with the view of Hare, described below, who compares doing conceptual analysis with describing the steps of a dance that one already knows well.

The *actual* current usage could be different from the *supposed* features of current usage as reported by philosophers and others after naive, untutored, attempts at introspection of their understanding of meanings. The incorrectness could be demonstrated by pointing at actual uses that are inconsistent with the introspective reports, or by showing that the supposed features if they existed would be inconsistent with aspects of usage that everyone agrees about, and which can be tested by observing what we say and do in our unreflective uses of the relevant words and phrases.

So the unreliability of many attempts to describe actual usage would be something like the unreliability of introspective reports by native speakers of the rules of the grammar they use, reports that could be checked against what they actually say when having real conversations. E.g. Kant wrote that every judgement is composed of subject plus predicate, whereas he wrote books full of sentences with far more complex and varied structures, e.g. conditionals, disjunctions, etc., refuting what he wrote, as does the proposition expressed by this sentence. So his claim, allegedly a priori, was, unusually for Kant, empirically false.

(In my [1978 book](#) I made a related distinction between 'common sense' and 'common opinion', not realising that I was probably echoing a distinction picked up from Ryle and others.)

Ryle clearly did not always assume, as some 'ordinary language' philosophers were supposed to, that current ordinary usage was always a standard of correctness. I presume he would accept that ordinary usage could be confused or even [inconsistent](#), which is why it may require philosophers to 'rectify' it.

From actual to possible usages (Added: 11 Feb 2007)

The wording in one of Tanney's sentences quoted above, points to the need for what I am calling 'logical topography'. She contrasts "the recipe theorist" who must devise recipes "by construing the role that ingredients, measures, and mixing and heating procedures play within a range of possible (successful) dishes" with the philosopher whose chart "deals with the various ways in which these concepts figure in the sayings (not only the describings) of people competent in their use."

The crucial difference is that the recipe theorist can devise new theories within the space of possibilities generated by ingredients, measures, mixing and heating procedures, whereas the philosopher is doing something closer to analysing and comparing some given set of recipes that are actually in use. The work of the former therefore has broader scope than the work of the latter. The recipe theorist who realises that the existing cooking technology supports a kind of recipe never before thought of is extending the set of recipes available for charting. But doing that requires an appreciation of the underlying 'logical topography' that makes the set of actual recipes, as well as alternative sets of recipes (e.g. used in other past present and future cooking cultures) possible, just as the topography of some terrain supports many different possible ways of dividing up that terrain (different geographies).

NOTE (added 20 Apr 2007):

I have discovered an earlier use of the phrase 'logical topography' near the beginning of this paper

Ansgar Beckermann, Property Physicalism, Reduction and Realization, in: M. Carrier and P. Machamer (Eds.) *Mindscapes. Philosophy, Science, and the Mind*. Konstanz: Universitätsverlag/ Pittsburgh: Pittsburgh University Press, 1997, pp 303--321.

Available at: http://www.uni-bielefeld.de/philosophie/personen/beckermann/prpph_ww.pdf

He writes:

In the course of this paper I will therefore do two things: At the start, I will try to sketch the logical topography of the "solution space" of the problem Kim is concerned with.

This appears to be used in the same sense of 'logical topography' as my usage, though it is not clear whether that was just a slip, and 'logical topography' was intended as 'logical geography', in Ryle's sense.

Is study of logical geography empirical linguistics?

It is tempting to interpret this study of 'logical geography' as a branch of empirical linguistics, as I believe the [Norwegian philosopher Arne Ness](#) did around that time: he tried to solve problems in conceptual analysis by doing surveys of linguistic usage, and opinions about usage.

However, Ryle and other analytical philosophers, studying logical geography, or doing conceptual analysis, thought (and many still think) that that is not an *empirical* investigation. Explaining why not is not a simple matter, and attempts to explain why are subject to the risk of claiming that the results of conceptual analysis are *analytic*, in something like the sense explained in ['Necessary', 'A Priori' and 'Analytic'](#) (1965), making the results of such investigations trivially true, i.e. true by definition.

This is closely related to [G.E.Moore's 'Paradox of analysis'](#). See also the Stanford Encyclopedia of Philosophy survey on [analysis](#).)

My 'personal tutor' at Oxford [Richard Hare](#) kindly gave me a lot of his time although I was not working in his research area. He was seriously concerned about the problem of clarifying what exactly philosophers were doing when they did conceptual analysis or investigated some area of logical geography. He appreciated the argument that analytical philosophers were merely doing armchair empirical research on the use of language, and attempted to refute that argument by distinguishing studying a language as an anthropologist does and studying the language as a native speaker does.

He sometimes described the process of conceptual analysis as "uncovering what you had learnt at your mother's knee" (or something like that). He seemed unconcerned that this appeared to shift the activity away from empirical linguistics at the cost of turning it into some sort of empirical psychology. He compared the process of doing conceptual analysis by trying out examples with the process of settling disputes about how steps in a well known dance go by actually *doing* the dance. See [Are Discoveries about the Uses of Words Empirical?](#) The Journal of Philosophy, November 7, 1957. This is almost the same paper as [Philosophical Discoveries](#). Mind, April, 1960.

I think Hare was right in claiming that there is an interesting difference between simply trying to remember how the steps go and actually getting into the swing of doing the dance to find out how the steps go. But it is not clear to me that what he wrote about that difference really explains why philosophers are doing something deep but not empirical.

If he had added that the theory of how the steps go might be articulated formally then tested for consistency, and then various consequences derived using logic or mathematics, then that would have shown that analytical philosophy can have much in common with deep theoretical sciences, which combine an empirical component with a significant mathematical or logical component which is not empirical, namely the study of relations between parts of the theory, and the study of implications of the theory. Compare Frege's work on arithmetical concepts mentioned below.

It is not clear to me whether any of the mid-20th century analytical philosophers ever satisfactorily clarified the nature of conceptual analysis or the study of logical geography, which I think amounts to the same thing.

Conceptual analysis as deep science (Updated 11 Feb 2007)

One fairly obvious option did not seem to tempt any of them, namely comparing conceptual analysis with the task of deep theoretical science, as suggested above. Doing deep science, e.g. newtonian mechanics, electromagnetic theory, quantum physics, molecular biology, theoretical linguistics (as opposed to the kind of shallow science that merely collects data and looks for correlations), involves coming up with a theory that cannot be derived from the evidence, but can be tested against it, where the theory has a collection of core 'axioms' embedded in a system of logic or mathematics in which inferences can be made.

This could be done for a theory about the structure of a collection of concepts. Then demonstrating that such a theory (a) is internally consistent, and (b) has the right set of consequences to explain facts of usage about which there are (usually, though not always) no dispute, would be more like a mathematical and therefore non-empirical task than like the empirical task of collecting evidence about usage using surveys.

Had the philosophers taken that line they could have said that doing logical geography was not empirical because it was not a study of relations between words and meanings, but rather a search for a collection of structures with certain logical properties, including a demonstration of the internal consistency of such a collection, and the logical consequences.

But that would still leave the empirical nature of the claim that the alleged properties matched actual linguistic usage, just as, after all the mathematical derivations regarding quantum mechanics, there remains an empirical question whether that theory fits the way the world works.

Such an approach to conceptual analysis would be very similar to the study of linguistic competence as defined by Noam Chomsky, in [Aspects of the theory of syntax](#) (1965), though that had not yet been written at the time *The Concept of Mind* was published.

However, many of the philosophers of that time did know about Frege's monumental work attempting to explain how our concepts of number theory work, and they could have used that as a model of philosophical theory construction that is neither trivial definition-writing nor empirical generalisation from observations. Not only was it not trivial -- it turned out to be inconsistent, as Russell showed! However, consistent subsets of Frege's theory were extremely important, and later work produced consistent variants of the theory, though not everyone was happy with the results.

For an overview of Frege's work in this area see the Stanford Encyclopedia of Philosophy article [Frege's Logic, Theorem, and Foundations for Arithmetic](#)

Note that one of his books was translated into English by one of the important analytical philosophers:

Die Grundlagen der Arithmetik: eine logisch-mathematische Untersuchung Ueber den Begriff der Zahl, Breslau: w. Koebner, 1884; translated by J. L. Austin as *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*, Oxford: Blackwell, revised edition, 1953.

A collection of Frege's papers was also translated and published as a book read by many analytical philosophers.

Frege, Gottlob, *Translations from the philosophical writings of Gottlob Frege*. Translated by Peter Geach, and Max Black. Oxford : Basil Blackwell, 1952.

This notion of a theoretical structure that can be precisely specified and from which many, or even infinitely many, conclusions can be drawn or constructions made, among which there are different possible ways of introducing subdivisions, provides the basis for the distinction I am trying to make between *Logical Topography* (the structure of the space within which the divisions can be made), and *Logical Geography* (a particular set of subdivisions that happens to be used in some culture).

We can distinguish

- a collection of empirical facts and empirical conjectures about some portion of reality,
- a theory T purporting to explain those facts (and possibly others),

(creation of such a theory T, given the facts, is a process of abduction, as described below -- however, alternative, rival, theories often result from abduction).

- the empirical conjecture that T is the correct, or the best available, explanation of some collection of empirical facts,
- a collection of non-empirical and *necessary* truths about:
 - the consistency of T
 - the consequences of T, either on its own or in combination with additional empirical

- propositions,
- A larger collection of non-empirical truths about:
 - the consistency of T combined with various other theories,
 - the relations between T and those theories (e.g. is it derivable from some of them, are some of them derivable from it)
 - the consequences of T combined with various other theories and empirical propositions.

Similar distinctions can be made about games defined by collections of rules.

E.g. if a group of people happen to play a game with the rules of chess (or rules isomorphic to the rules of chess) that is an *empirical* fact. But whether anyone actually plays the game or not, it is a *necessary* consequence of the rules that certain board positions can occur in a game and others cannot. Moreover, it is a necessary consequence that there are very many distinct kinds of board positions that can occur (not infinitely many, but more than the number of electrons in the universe, I believe -- I have not done the calculations).

There are other necessary truths about the consequences of modifying the rules in various ways, and truths about what possible sequences of moves remain after a certain board position has been reached.

The set of possible (interesting or uninteresting!) two-player board games using a square or rectangular grid with a set of pieces on the board that can be moved and or removed from or added to the board, defines a (large) logical topography. The logical geography of actual games of that sort played in some culture uses a small subset of the possibilities defined by the topography.

The task of starting from an explicit theory and working out consequences of the theory, on its own or in combination with other theories and empirical propositions, is not an empirical task (though there are variants of it that are empirical, e.g. finding out what the actual consequences are in certain circumstances where other unknown factors operate).

A harder non-empirical task: Abduction of a new theory

However, besides those non-empirical tasks, there is a much harder (abductive) task which is also non-empirical, namely, given some collection of facts or supposed facts F, and some collection of theories T, find a new theory T' which

- is internally consistent
- is consistent with F and T
- is capable, in conjunction with T, of explaining the collection F as being necessary consequences of T and T', though T alone does not explain the collection.

The term 'abduction' was invented by C.S. Peirce, as far as I know.

NOTE:

Abduction is sometimes presented as a third form of reasoning, of which the first two are deduction (derivation of consequences that necessarily follow from some assumptions), induction (derivation of consequences that are more likely to be true than false if the assumptions are true).

However, this is very misleading because abduction, at least as defined here, can lead to the creation of several rival explanatory theories, starting from the same set of assumed facts and theories. The mere fact that such a theory T' has been created which satisfies the conditions given is not the slightest reason for regarding it as true, even if F and T are true. The most that can be claimed, in some cases, is that T' is better than proposed rival theories which explain the same phenomena. However, as K. Popper, and later his pupil, Imre Lakatos, pointed out, it is not in general possible to determine in any decisive way which of two theories is better at a

particular point in time, e.g. because available technology is inadequate for the task. It can take a long time before it is clear that, for example one theory is the core of a progressive research programme, while the rival research programme is degenerating, in Lakatos' terminology. However it is always possible for the tables to turn later on, in the light of new empirical discoveries, as happened to the particle theory of light which appeared for many years to have been defeated by the wave theory of light, following the discovery of interference phenomena.

(There is no universally agreed terminology relating to these topics: it's another area whose logical topography has probably not been adequately studied, and which allows different logical geographies to be constructed. For example, for consistency with the above definitions, what is now called 'Inductive Logic Programming' (ILP) should probably have been called 'Abductive logic programming', insofar as ILP generates new explanatory hypotheses given some data.)

NOTE:

For a recent discussion of Abduction in the context of philosophy of science see Alan Musgrave, [The Miracle Argument For Scientific Realism](#), in *the Rutherford Journal* Volume 2, 2006-2007.

Musgrave notes that some philosophers have argued that if abduction leads to a theory like T' discussed above, then there is reason to believe it is true, but fails to point out that abduction can produce many mutually inconsistent theories from the same starting point, so if it provides reason to believe any of them is true it provides reason to believe all are true. So in itself it gives no help in selecting good from bad theories. For that, more than abduction is required.

It is surprising that Musgrave never mentions the work of Lakatos on 'progressive' and 'degenerating' research programmes, since they co-edited the book in which Lakatos published his first famous paper on this topic.

[Falsification and the Methodology of Scientific Research Programmes](#)

NOTE: (Added 1 Oct 2007)

The previous remarks would lead most readers to assume that abduction involves the use of sentences in some logical language, i.e. a language in which it is possible to perform deductive (and possibly also inductive) reasoning.

However it is clear that much causal reasoning in everyday life, in science, and in engineering, and also much mathematical reasoning, is not restricted to the use of linguistic and logical formalisms but also includes various kinds of spatial, geometric or diagrammatic reasoning. In this 1971 paper

[Interactions between Philosophy and Artificial Intelligence: The role of intuition and non-logical reasoning in intelligence.](#)

I argued that not only Fregean but also analogical forms of representation can be used in valid, even rigorous reasoning. If that is correct then it is to be expected that some forms of abductive reasoning, instead of producing logical or verbal theories, will use spatial structures or working models to express hypotheses. Then the phenomena to be explained will be derived by geometrical reasoning, or running simulations, rather than by logical deduction.

Related recent work includes these presentations on Kantian and Humean concepts of causation, and the perception of spatial processes.

- [Two views of child as scientist: Humean and Kantian \(PDF\)](#)
- [Work with J. Chappell on causation in animals and robots.](#)
- [A \(Possibly\) New Theory of Vision \(PDF\)](#)

It is not unknown for philosophers to attempt to present their theories using diagrams. I think it is rare for the diagrams to have sufficiently precise semantics to be rich in explanatory power, in contrast with diagrammatic formulations in science, engineering and mathematics.

NOTE: (Added 1 Oct 2007)

I have just discovered that the philosopher Lorenzo Magnani has written some related things, though they differ in many details. Further information is available here

http://www.unipv.it/webphilos_lab/cpl/

It is fairly obvious that many great advances in science have involved an abductive process of the sort described here, often producing competing new theories which then have to be tested over time, and often introducing new concepts, that are not definable in terms of the previously used concepts.

I believe that much deep philosophy can also be construed as offering a theory T' satisfying such conditions. Examples would include some of Kant's philosophy and Frege's philosophy of mathematics, much philosophy of science, and, of course, what Ryle was trying to do in *The Concept of Mind*, though he was using conceptual tools that were inadequate for the task. (Most of my own work is also of this form -- and blurs boundaries between philosophy, psychology and biology.)

Once such a theory T' has been developed, then, together with previous theories T, and other facts about the world, it will generate consequences. If the new theory T' is deep and general it will generate a lot of different sorts of consequences about things that are possible, e.g. states, events, and processes that can occur, as the set of rules for a game like chess or go generate consequences about possible board positions and possible sequences of moves. Some of those consequences may be surprising.

They need not be testable consequences if they are of the form 'so and so can occur', since no observation can ever refute such a proposition, though an observation of a so and so can prove it true. (This point was part of an argument against Popper's philosophy of science in [Chapter 2](#) of *The computer revolution in philosophy*.)

That collection of possibilities defines a 'logical topography' which admits various ways of dividing up the possibilities, e.g. classifying board positions, classifying strategies, classifying types of games, classifying types of reasoning or cogitating.

For each division there will be a collection of concepts and a set of relations between those concepts, which we can think of as constituting a logical geography, made possible by the underlying logical topography, which also makes possible alternative logical geographies. Some of the concepts will be quickly learnt by novices, whereas others may only become evident after years of study by experts.

As far as I can tell, neither Ryle nor Hare noticed the possibility of making this sort of distinction, between (a) collecting what are essentially empirical facts about how concepts are actually used in some community (even if they are not discovered or checked using standard methods of empirical science) and (b) doing the non-empirical work of both analysing the structures involved in various logical topographies and logical geographies, and searching (abductively) for a new explanatory theory that is capable of accounting for the empirical facts about usage.

Besides the empirical questions about which sets of concepts are used in a particular culture, there are also empirical questions about which theories using those and other concepts are good scientific theories explaining the relevant empirical facts.

I think [Peter Strawson](#) came close to this with his notion of 'Descriptive metaphysics', contrasted with 'Revisionary metaphysics' in his book [Individuals: An Essay in Descriptive Metaphysics](#). London: Methuen, 1959. As far as I know, however, he did not take the step of allowing that the outcome of a study of descriptive metaphysics (another name, I think, for what Ryle called 'logical geography, though perhaps on a grander scale) could include parts that differ from one culture to another. Had he done so that would have made clear the empirical component of the investigation. Whether he is right in claiming that there is a core that cannot vary across cultures is a separate question. One way of demonstrating him to be wrong would be to produce kinds of intelligent robots using conceptual systems without the particular core that Strawson claimed was indispensable.

(Brian Cantwell Smith's book [On the origin of Objects](#) (1996) attempts to address this question.)

I think Strawson would claim that we would not be able to communicate with such robots, any more than we can communicate with earthworms. But maybe we could learn to do so by absorbing their conceptual system, as anthropologists learn to communicate with alien cultures. I don't think he demonstrated the impossibility of such a scenario, but that's a topic for another occasion. (Also much discussed by Quine in his famous 'gavagai' example.)

Wittgenstein's claim that he was doing anthropology when doing philosophy may have expressed a view that is related to the claim that all philosophical conceptual analysis has an empirical content, something of the form: '*this set of concepts is used by this community, or this animal species (humans)*'. But that may be too generous to Wittgenstein.

The distinction between what I am calling 'logical topography' and 'logical geography' is very nearly made in Tanney's example of all the uncharted "implication threads" in the recipes used by an expert chef.

Another kind of logical geography? Logical topography.

For some time I thought Ryle used the label 'Logical Geography' to refer to what I am calling 'Logical Topography', the study of which I think is a more important task, and much harder and deeper than merely charting the actual patterns of usage and their implications. Moreover it is a prerequisite for any 'rectification' of the kind [Ryle was hoping for](#) when he wrote the introduction to *The Concept of Mind*.

Analysing logical topography involves

not just exploring the connections in *existing* usage, but exploring the variety of *possible* concepts that can be defined in connection with the subject matter that is being investigated, analysing the relative merits of different partitionings, and showing how actual concepts do or do not relate to those.

The latter is what I have been trying to do for about three decades, most recently in connection with discussion of '[architecture-based conceptions of mind](#)'. Investigation of logical geography in this second sense (i.e. studying logical topography) is a prerequisite for doing what P.F. Strawson labelled 'Revisionary Metaphysics' in [Individuals](#) (1959).

The linguist John Lyons in his two volume book *Semantics* points out that names for floor coverings (e.g. 'rug', 'carpet', 'doormat', 'bathmat', etc.) as used in different languages refer to different ways of sub-dividing possible coverings. So a linguist studying usage in English and in French might expose two logical geographies. But in order to talk about them he would have to presuppose a logical topography, a space of possibilities generated by the nature of floors, and coverings and possible patterns of use of floors and their coverings.

The difference between logical geography and logical topography is a bit like the difference between describing the *actual* division of some area of terrain into political regions, and describing the features of the terrain that would usefully support different *possible* ways in which the terrain could be divided into political regions. That analogy is what led me to the phrase 'Logical Topography'. (Not all studies of logical topography would be concerned with a continuum of possibilities, however).

An example from Physics

An example I often give is the way creation of the theory of the atomic structure of matter allowed different ways of partitioning kinds of stuff, which, for the purposes of science and engineering, improved on the 'logical geography' of kinds of stuff that existed previously in ordinary language. A very useful partitioning is illustrated in the periodic table of the elements, and elaborated by the theories of how different chemical substances can be formed by different combinations of the component atoms. That's an example of a discrete topography. An example of the way in which atomic theory advanced previous concepts was the ability to distinguish isotopes, so that we could, for example, talk about carbon12 and carbon14, instead of arguing about which is 'really' carbon.

An example from software engineering

Another example would be taking a specification of some information processing architecture, then analysing the varieties of states and processes that can occur in that architecture and investigating alternative ways of clustering them. E.g. merging sets of states in attractor basins into equivalence classes would be one way. A different way of dividing states or processes might be in terms of whether some subset of the architecture happens to be active or not, or whether it is or is not connected with other subsets. (This might, for example, be part of a new logical geography mapping out the terrain loosely indicated by ordinary words like 'notice', 'attend', 'conscious', 'experience'.) There are very many other ways of classifying states and processes in a complex system made of many interacting components. Maintenance engineers and operators concerned with some complex control system will often work out what the interestingly different states and processes are, from the point of view of facilitating their jobs, e.g. states that are associated with particular kinds of software bugs, or states that are potentially dangerous, or states in which a particular positive feedback loop is active.

The divisions that may be of interest to one lot of users will not necessarily be the same as the divisions of interest to the others. For instance, one lot may be concerned with different states and processes that are relevant to finding out which component is faulty when the system goes wrong. Another may be more interested in the different contexts of use, and the different functions, of the system when everything is working properly. (It is often the case that maintenance engineers and users of a complex system use different collections of concepts in thinking about the system.)

Likewise divisions of states and processes made possible by the workings of a human mental architecture that are useful for normal people, including garden-gate gossips, may be different from partitions that are useful for psychiatrists. Novelists could be somewhere in between.

(Some analytical philosophers get very upset by the notion that investigation of an architecture-based logical topography can lead us to improve on our pre-theoretical concepts of mind. I often wonder what Ryle would have thought, had he been exposed to these ideas, though not in the naive versions often described as 'functionalism', criticised [here](#) and later in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/vm-functionalism.html>)

Another example from software engineering

The many and various uses of the word 'ontology' in Software Engineering and Artificial Intelligence form a conceptual morass. Later I shall insert some notes here showing how to unravel the tangle of ideas, separating out philosophical issues related to descriptive metaphysics from other ideas related to program code, particular tools and algorithms, and design methodologies.

Weeding and cleaning out the tangled mess will require studying the logical topography of a host of problems connected with attempts to design more or less human-like machines.

Arguing about which definition is right, or best, is usually pointless

The fact that an information processing architecture generates a space of possible states, events and processes that can be carved up in different ways for different purposes, provides an opportunity for a deep and general study of that space, and the implications of different ways of carving it up. From that perspective there is no point in arguing about which states and processes *should* be describable as "emotional" or "having a motive". However there is much interest in trying to understand the 'logical topography' of possible sets of states and processes associated with different architectures, in a way that relates to a deep explanatory theory.

If we can work out what sorts of distinctions *can* be made in the framework of that topography, and what the costs and benefits of choosing between alternative ways of carving up the terrain are (the trade-offs) then that can be a real advance in knowledge. Selecting one or other taxonomy can then be done on the basis of that new knowledge.

For example, biologists used to divide plants and animals into categories based on their visible features (including habitats and behaviours). That is still a possible way to do it, but we now understand that the logical topography allows alternative ways, including defining species in terms of evolutionary history, or genetic makeup. And because that option is associated with a deep explanatory theory we now choose to class whales and dolphins as mammals rather than as fish.

Problems of 'cluster concepts'

Many of our concepts, especially pre-scientific concepts, are cluster concepts. (I think the term was originally coined by D. Gasking.) If C is a cluster concept, then there's a collection of different criteria involved in the the application of C to instances, but no subsets of those criteria definitely form either necessary or sufficient conditions. So there may be some instances for which some of the criteria suggest that C applies whereas other criteria suggest that C does not apply to those instances. If C is a cluster concept there is *no* right answer to the question: are these things *really* of type C?

Sometimes other labels are used for this or closely related notions, e.g. concepts with "open texture" (Waismann), "family resemblance concepts" (Wittgenstein), "[essentially contested concepts](#)" (Gallie), and no doubt others. (This is not the same as simply having an indeterminate boundary in a range, e.g. the boundary between red and orange.)

When we already use concepts with that sort of indeterminacy then if two rival refinements are proposed that make the conceptual boundaries more determinate, there is often no answer to the question whether one of the new definitions is *correct*, as an attempt to clarify or improve the *old* concepts, though we can argue about which revision is *better* for some purpose, as illustrated below in connection with mathematics and wheel-making.

Note added 24 Feb 2016

I have recently realised that many of the concepts that I had regarded as "Cluster concepts" could best be regarded as concepts exhibiting "parametric polymorphism" as discussed in <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/family-resemblance-vs-polymorphism.html>

Problems of inconsistent concepts (added 22 Oct 2007)

There are inconsistent concepts connected with logical and semantic paradoxes, such as the concept of *the set of all sets that do not contain themselves* (Russell's paradox) and the concept of a heterological word or concept (e.g. a word is heterological if and only if it does not describe itself).

However, it is arguable that there are many other concepts that are inconsistent in other ways, including for example, some people would say, the concept of an all powerful, all knowing, perfectly good God, or certain concepts of free will. I have argued that the ordinary use of the word 'consciousness' can be shown to be inconsistent, for example because people will say that in order to have an emotion you need to be conscious, fear is an emotion, dreams occur during sleep, being asleep involves being unconscious, and it is possible to experience fear when you dream you are being chased by a hungry lion.

For more on the problems of trying to use ordinary notions of 'consciousness' to define a subject of scientific investigation see

<http://www.cs.bham.ac.uk/research/projects/cogaff/03.html#jcs03>

Aaron Sloman and Ron Chrisley,

Virtual Machines and Consciousness

Journal of Consciousness Studies, 10, No. 4-5, 2003.

In [Varieties of Atheism: What is analytical atheism?](#), I have offered a partial analysis of the logical topography of the space of phenomena referred to as theism, atheism, agnosticism, pointing out that there is a concept of "analytical atheism" supported by that logical topography, whose importance is rarely noticed. That document also includes a brief tutorial on the use of conceptual analysis to expose disguised nonsense for what it really is.

In various papers and presentations, e.g. [here](#), I have tried to show that various muddles generated by trying to analyse many familiar concepts of mind, e.g. "emotion", "mood", "attitude", "desire", "preference", "value", and others point to a need to use the notion of an *information-processing architecture* to generate the logical topography that we are trying to parcel up by defining those concepts. That is because many of our ordinary concepts of mind are best construed as ["architecture-based" concepts](#).

Do dynamical systems with attractors have motives?

In a paper on [Requirements for a Fully Deliberative Architecture](#), written earlier this year, I made the following comments:

This is not a debate about definitions.

Note that I am not claiming that my use of the word 'deliberative' is right and others wrong: it is almost always silly to argue about the 'correct' definition of any word in the context of scientific debates. What is important is to get an overview of the range of phenomena that we are trying to understand and to find ways of dividing them up that lead to deep explanatory theories. (Unlike: earth, air, fire and water, or Lewis Carroll's walrus' list: 'shoes and ships and sealing-wax and cabbages and kings'.) The philosopher Gilbert Ryle used to describe this as exploring the 'logical geography' of a set of concepts. Analysing logical geography is far more important than arguing about definitions. (For a recent

introductory overview of Ryle's ideas see Julia Tanney's paper [Rethinking Ryle.](#))

It is perhaps surprising that nobody has presented me with an argument that water is deliberative because if a stream of water starts flowing down a gully and comes to a fork in the gully, where one branch soon meets a dead end (where ground rises) whereas the other branch goes on down the hill, the water will 'try' both branches, and most of it will select the better branch to flow down. Perhaps cases like that explain why Aristotle thought water had the goal of getting to the centre of the earth?

NOTE

At the time I wrote that, I had forgotten that George Kiss had suggested that motives should be thought of as attractors in dynamical systems, about 30 years ago.

Little did I know that a few months after I had written the above, Simon McGregor (a PhD student at Sussex University) would write to me after I had made statements about possible minimal requirements for a system to have motivational states, in an email discussion:

"I'm not yet convinced that there is a consistent theory of what constitutes motivation and intentional action which excludes falling rocks. It may turn out that the best way to see these things has the logical consequence of requiring us to recognise a trivial version of motivation in rocks."

(Simon McGregor, personal communication.)

My response is that what's consistent depends on what else is in the theory, and what's best may depend on what the purpose of having the theory is.

For instance, it would be possible for mathematicians to decide not to call circles ellipses, and the result might be consistent terminology, but with a lot of awkwardness. E.g. many theorems might need 'and circles' or 'or circles' to be added. In contrast, for the purposes of a wheel-maker training apprentice wheel-makers it may be very useful to stress the importance of not making elliptical wheels.

So different logical geographies superimposed on the same logical topography are useful in mathematicians' and in wheelwrights' workshops where even '[round](#)' is taken to imply non-elliptical.

If it turns out to be useful to extend the concept of having a motive to include any dynamical system that is in an attractor basin then what I have been assuming are necessary architectural requirements for having a motive are not necessary.

Whether that extension is useful or not, we can still learn by discussing differences between that case and others.

The issue was discussed in a fairly abstract way in

A. Sloman, R.L. Chrisley, M. Scheutz, The architectural basis of affective states and processes, in *Who Needs Emotions?: The Brain Meets the Robot*, Eds. M. Arbib and J-M. Fellous, Oxford University Press, pp. 203--244, 2005.

(Online [here.](#))

My general view about these matters is that the definitions are not worth arguing about except in the context of a discussion of which of two or more rival *theories* is better.

If the discussion does not bear on such a decision, then which definition is best will be a matter of convenience, and what is convenient can vary with context. It is probably not convenient for a wheelmaker to abandon the notion that 'circular' and 'elliptical' are inconsistent descriptions, even though mathematicians think otherwise. If necessary, we can just use slightly different labels for the two concepts (like the labels for isotopes of a chemical element), or allow the context to decide which is in use: e.g. whether it's a mathematical discussion or a tutorial for apprentice wheel-makers.

(However, if you are trying to propose some new terminology for use in a well-specified scientific or engineering community then it may be worth trying to agree on something, after analysing the tradeoffs between the different options).

One reason why it is usually not worth arguing about definitions is that there is usually a deeper and more interesting task, which is exploring the logical topography of a more fundamental set of concepts, in the sense defined above.

Some ways of reaching agreement

People who are inclined to use pre-existing terms (e.g. 'motivation', 'emotion', 'consciousness', 'learning', 'intelligence') in different ways, can still agree, if they have a suitable explanatory theory, that there are significant differences in the phenomena they are discussing, e.g. differences between what happens when an unstable rock starts rolling downhill if tapped, and what happens when a woodlouse rolls itself into a ball when tapped, even if one of them wants to extend the notion of 'motivation' to both, and the other wants to include only the woodlouse.

So pointless disagreements can be replaced by collaborative, and possibly fruitful, analysis, if people adopt the right analytical stance.

Of course some people will use even more restrictive criteria and exclude the woodlouse. E.g. they may require the architecture to support explicit representations of desired states in advance of their realisation, which I suspect a woodlouse architecture doesn't.

Whereas a *Portia* spider does:

The *portia* spider works out a route to its prey then follows it even when it can no longer see the prey, making detours if necessary and avoiding branches that would not lead to the prey.

"By visual inspection, they can select, before setting out, which detour routes do and do not lead to prey, and successfully perform a detour with no further visual contact with the prey".

M. Tarsitano, 2006,

Route selection by a jumping spider (*Portia labiata*) during the locomotory phase of a detour,

Animal Behaviour, 72, Issue 6, pp. 1437--1442,

<http://dx.doi.org/10.1016/j.anbehav.2006.05.007>

The same sort of distinction might be made between different things that we are inclined to call 'desires' in humans, some of which merely involve powerful tendencies to act or decide in certain ways, while others involve an anticipated result.

(Added: 26 May 2009): The former could be examples of 'architecture-based' as opposed to 'reward-based' motivation, as explained [here](#).

Interesting differences concern whether behaviour can be changed by learning, the ability to reverse actions, the variety of ways in which processes can be triggered and prevented, the different sorts of consequences of the behaviour, the variety of means by which the motivation can be realised, the variety of ways of dealing with obstacles, whether multiple motivations can coexist, how conflicts are resolved, etc. We could introduce many subdivisions -- which may or may not be worth explicitly labelling, depending on why we are doing it.

Delving into the possible explanatory conceptual structures inherent in some physical, geographical, social, computational or psychological system is often both fun and very challenging. It can also contribute to science and engineering. Perhaps everybody should learn how to do it: it does not come naturally to most people.

There is a short introductory tutorial on conceptual analysis in this discussion of varieties of atheism: <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/varieties-of-atheism.html>

A more demanding tutorial is in Chapter 4 of *The Computer Revolution in Philosophy* <http://www.cs.bham.ac.uk/research/projects/cogaff/crp/#chap4>

Logical topography of Meta-Requirements for Robots

A draft document on the analysis of notions like 'robust', 'flexible', 'creative', 'intelligent', 'autonomous' in the context of specifying requirements for future robots can be seen as an exercise in Logical Topography as defined here. The document is [A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots](#)

Logical topography of concepts of deliberative agents

Another example of analysis of logical topography is a discussion paper on varieties of concepts of 'deliberative' capability. [Requirements for a Fully Deliberative Architecture \(Or component of an architecture\)](#)

Some earlier work on logical topography of 'better', 'good', 'ought', 'identity' etc.

Added: 25 Feb 2007

I have just realised that some of my earlier work in meta-ethics can be seen as exploring logical topography and using abduction to propose an explanatory theory, namely:

["How to Derive 'Better' From 'Is'"](#) in *American Philosophical Quarterly* 6, 1969

["'Ought' and 'Better'" in *Mind* Vol 79, 1970.](#)

These are contributions to what philosophers call '[meta-ethics](#)', as well as proposing what I believe was a new approach to a subset of linguistic theory treating words like 'better' and 'ought' as higher-order predicates some of whose arguments tend to be suppressed. To see the context, look at Jimmy Lenman's very useful annotated [Bibliography of meta-ethics](#)

(Added 25 Feb 2007)

A related paper challenged some assumptions about how the concept of personal identity relates to ethical questions, by presenting an example of a theoretically possible future technological advance that would pose questions about identity whose answers would depend on social and ethical views rather than the existence of some necessarily 'right' answer:

[New Bodies for Sick Persons: Personal Identity Without Physical Continuity \(PDF\)](#) Analysis vol 32 NO 2, December 1971, pages 52--55

This is a topic that goes back to ancient greek philosophers, e.g. [Heraclitus](#): 'You cannot step into the same river twice'. There is also the story of [the ship of Theseus](#) which was continually being rebuilt while in use, and the borrowed axe, whose handle broke and was replaced while it was on loan, and then whose blade broke and was replaced: did the lender return the borrowed axe? Exploring the logical topography of 'identity' and related concepts shows that there are alternative possible logical geographies, and there is no right answer to such questions. A full analysis of the logical topography of notions related to 'identity' would probably need to describe the biological evolution of forms of representation of various kinds of entity in the environment, and how those forms changed under pressure of changing requirements, including, eventually, social, cultural and ethical requirements. In some cultures that probably included pressures from religion, whose myths and inspirational ideas depended on pushing the notion of personal identity in a particular direction, which could be opposed by other viewpoints.

The issues also arise in connection with problems of designing intelligent robots and other machines. Here the discussion could relate the benefits of using alternative notions of identity to different practical applications of such machines.

NOTE:

After writing the first draft of this paper, I discovered an excellent online essay by Dennett. 'Re-Introducing The Concept of Mind' available online in [PDF](#) and [HTML](#) formats.

NOTE Added 21 Apr 2008:

I have just discovered an interesting and relevant article by Jerry Fodor

<http://www.lrb.co.uk/v26/n20/print/fodo01.html>

Water's water everywhere

Jerry Fodor

Review of *Kripke: Names, Necessity and Identity* by Christopher Hughes

In London Review of Books 21 October 2004

[I may add further comments on this, and how it relates to the distinction between logical geography and logical topography later.]

Note Added 1 Jan 2010:

[I found this link with relevant comments](#) (click on the upper [+/-])
by David Jhave Johnston (an artist-poet with a computational background) on a mostly artistic
web-site
<http://www.year01.com/html/BAndH.html> located here <http://www.year01.com/>

His Masters thesis is here

<http://www.year01.com/jhave/thesis/FINAL/>
FERTILE SYNTHESIS: EMOTION IN ONLINE DIGITAL POETRY
by William David (Jhave) Johnston
Thesis submitted in partial fulfillment of the requirements for the degree of Master of
Applied Science
School of Interactive Arts & Technology, Simon Fraser University, Summer 2007

His homesite is <http://glia.ca>.

His work (which I have not yet studied closely) draws attention to the important point that one of the things artists do is draw attention to some of the structure in both natural and artificial dynamic topographies, in addition to creating new ones. Footnote 6 on Page 3 says: *"What I am suggesting differs from Sloman's concept in advancing a notion of the terrain as a dynamic, and drawing an analogy between truth-validity and survival."*

I agree with him that there are strong connections between truth and survival -- though biological evolution shows how errors can also survive.

We are both interested in dynamic topographies, though possibly on different time-scales. If anything I wrote above suggested that all logical topographies are static (as the periodic table of the elements is) then I did not express my intentions clearly. In particular, a functioning information-processing architecture, as sketched in the work on [the CogAff project](#), is highly dynamic and the sub-states and sub-processes present complex and fluid logical topographies which our ordinary labels for mental states and processes (believe, desire, intend, emotion, understand, prefer) carve up into logical geographies that have proved useful for the purposes of everyday communication but which are grossly inadequate for the purposes of scientific understanding.

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