Mental Maps[1]

Ben Blumson
National University of Singapore

Abstract

It's often hypothesized that the structure of mental representation is map-like rather than language-like. The possibility arises as a counterexample to the argument from the best explanation of productivity and systematicity to the language of thought hypothesis – the hypothesis that mental structure is compositional and recursive. In this paper, I argue that the analogy with maps does not undermine the argument, because maps and language have the same kind of compositional and recursive structure.

I

Belief, according to many, is a map by which we steer (Ramsey, 1931, 146; Armstrong, 1973, Lewis, 1994, 310-311; Braddon-Mitchell and Jackson, 1996, 177-184). An important place in which this analogy with maps arises is in evaluating the language of thought hypothesis – the hypothesis that the structure of mental representation is language-like (Lewis, 1994, 310-311; Braddon-Mitchell and Jackson, 1996, 177-184; Camp, 2007). One consideration supporting the language of thought hypothesis is an argument from the best explanation: the hypothesis that thought is structured like a

[1] Thanks to David Bourget, David Braddon-Mitchell, Christopher Gauker, Frank Jackson, Raamy Majeed, Luca Moretti, Michael Rescorla, Brad Richards, Daniel Stoljar, Weng-Hong Tang as well as audiences at the University of Sydney, the National University of Singapore, the Australian National University and the 2007 meeting of the Australasian Association of Philosophy in Armidale.

language, according to the argument, is the best explanation of the productivity and systematicity of mental representation (Fodor, 1975; 1987; 2008).

The analogy with maps arises as a counterexample to this argument from the best explanation: the hypothesis that mental structure is map-like is purported to provide at least as good an explanation of the mind's systematicity and productivity as the hypothesis that mental structure is language-like. Since the hypothesis that mental structure is map-like, according to its proponents, can explain everything the language of thought hypothesis can, neither hypothesis is supported by an argument to the best explanation (Braddon-Mitchell & Jackson, 1996, 177-184). This paper argues – to the contrary – that the analogy with maps does not undermine the argument from the best explanation of productivity and systematicity to the language of thought hypothesis.

The standard way of supporting the argument for the language of thought hypothesis against this objection is to argue that the language of thought hypothesis explains some things, apart from systematicity and productivity, which the analogy with maps cannot. It's sometimes argued, for example, that the analogy with language can explain the ability to reason whereas the analogy with maps cannot, because there are inferential connections between sentences but not between maps (Bermudez, 2003, 160-2; Devitt, 2006, 146-7). However, I will argue that maps and other kinds of pictorial representation have logical structure of the same kind as sentences, so the analogy with maps explains everything the analogy with language can.

Instead of supporting the language of thought hypothesis by arguing that the analogy with maps cannot explain everything the analogy with language can, I will argue that the hypothesis that thought is map-like does not offer a truly alternative explanation of the systematicity and productivity of thought. The analogies with maps and language, I will argue, both offer the same explanation of the systematicity and productivity of thought, because both suggest that mental structure is compositional and recursive. The hypothesis that mental structure is map-like is not a genuine alternative to the hypothesis that it is language-like, because maps and language both share the same kind of structure.

Despite agreeing that maps differ from sentences in virtue of their resemblance to what they represent, I will also argue that maps do not differ from sentences in possessing a different type of compositional and recursive structure. Just as whether or not a language is mediated by arbitrary conventions or psychological necessity does not effect the syntax and semantics of that language, the fact that map-like representation is mediated by resemblance, I will argue, does not affect the question of whether or in what way maps are compositional and recursive. Maps suggest that thought is mediated by resemblance (O'Brien & Opie, 2004), but not that thought differs from language in structure.

Two clarifications. First, the analogy between maps and thought is extremely suggestive. Nevertheless, this essay is exclusively concerned with how that analogy bears on the arguments from the best explanation for the language of thought hypothesis, rather than with connected debates between, for example, possible worlds or structured content, connectionism or classicism and semantic holism or atomism (all discussed in Lewis, 1994). Whatever other disanalogies they might suggest between mental and linguistic representation, the analogy with maps and the analogy with language both support the same explanations of systematicity and productivity.

Second, I shall follow Fodor (1975; 1987) and Braddon-Mitchell and Jackson (1996) in treating the argument from systematicity and productivity as an abductive argument. There are also deductive versions of the argument, according to which compositionality and recursivity is entailed by, rather than merely explanatory of, systematicity and productivity (Davies, 1992; Lycan, 1993; Rey, 1995). The analogy with maps arises here not as an alternative explanation of the productivity and systematicity of thought, but merely as a counterexample to glossing the conclusion of the argument from productivity and systematicity as establishing that thought is language-like.

II

The language of thought hypothesis is purported to be the best explanation of the productivity and systematicity of thought. Systematicity is the connection of the ability to produce or understand some representations with the ability to produce or understand others (Fodor, 1987, 149). Thought is systematic because the ability to produce some thoughts is connected to the ability to produce others: the ability to think that John loves Mary, for example, implies the ability to think that Mary loves John (Fodor, 1987, 149-51).[2] Productivity is the ability to produce or understand an infinite number of representations. Thought is productive because the thoughts I have had are just a small fraction of the infinitely many thoughts I'm able to have (Fodor, 1987, 147-8).

Prima facie, both the analogy with maps and the analogy with language are equally well placed to explain the productivity and systematicity of thought. Language is systematic because the ability to understand some sentences is connected to the ability to understand others: anyone with the ability to understand the sentence 'John loves Mary', for example, also has the ability to understand the sentence 'Mary loves John' (Fodor, 1987, 149-151). And language is productive because people have the ability to produce and understand an infinite number of novel sentences in their language, even though they have only heard a finite number of them (Fodor, 1987, 147-8).

Likewise, maps are systematic and productive. Maps are systematic because the ability to understand some maps is connected to the ability to understand others: anyone who has the ability to understand maps which represent that Collingwood is east of Fitzroy, for example, also has the ability to understand maps which represent that Fitzroy is east of Collingwood. And maps are productive because people have the ability to understand an infinite number of maps, even though they have only previously encountered a finite number of them (Braddon-Mitchell & Jackson, 1996, 182). So, prima facie, the analogies

[2] Opinion divides on whether the systematicity of thought is *a priori* (Evans, 1980, 100-104; Davies, 1992) or *a posteriori* (Fodor, 1975; 1987): this is relevant to the question, not raised here, of whether the argument from systematicity is abductive or deductive.

with maps and language are equally well placed to explain the productivity and systematicity of thought.

Note that systematicity doesn't require that one is able to produce or understand all representations recombined from parts of representations one is able to produce or understand. In particular, the ability to produce or understand nonsense is not required: the ability to entertain thoughts about green paint and innate ideas, for example, does not entail the ability to entertain thoughts about green ideas or innate paint, because thoughts about green ideas and innate paint do not make sense. Similarly, although one cannot understand all maps recomposed of arbitrary parts of maps one can understand, this doesn't show maps aren't as systematic as language, since recombinations of arbitrary parts of maps don't always make sense.

III

Languages are compositional in that the meaning of whole sentences depends on the meanings and arrangement of their parts. The meaning of 'Theaetetus flies', for example, depends on the meaning of 'Theaetetus' the meaning of 'flies', and the order in which they are concatenated. This suggests, according to David Lewis, that "Mental representation is language-like to the extent that parts of the content are the content of parts of the representation [and that] If our beliefs are "a map ... by which we steer", as Ramsey said, then they are to that extent not language-like" (Lewis, 1994, 310-311). So the language of thought hypothesis, according to Lewis' suggestion, is the hypothesis that parts of what thoughts represent are represented by parts of those thoughts.

But this characterisation of the hypothesis that thought is language-like doesn't distinguish the explanation it provides from that provided by the analogy with maps. Just as parts of sentences represent parts of what whole sentences represent, parts of maps represent parts of what whole maps represent. Parts of the map of the Earth, for example, represent parts of the Earth: the southern hemisphere is represented by the map's lower half and the northern hemisphere by the map's upper half. If the language of thought

hypothesis is merely the hypothesis that parts of what thoughts represent are represented by parts of those thoughts, then the analogy with maps offers the same explanation of the systematicity and productivity of thought as the analogy with language does.

Moreover the thesis that parts of thoughts represent parts of what whole thoughts represent is not sufficient to explain productivity and systematicity. Take, for example, the French flag. The red, white and blue parts of the flag represent the three parts of what the whole flag represents: liberty, equality and fraternity. But the system to which the French flag belongs is neither systematic nor productive, since the ability to understand the French flag isn't connected to the ability to understand any other flag. So the characterization is too weak for two reasons: it neither distinguishes the explanation offered by the analogy with language from that offered by the analogy with maps nor explains the phenomena of productivity and systematicity.

IV

This suggests that a stronger characterization of the language of thought hypothesis is required both to explain systematicity and productivity and to distinguish it from the hypothesis that thought is map-like. As Jerry Fodor has suggested "the LOT story amounts to the claims that (1) some mental formulas have mental formulas as parts; and (2) those parts are "transportable": the same parts can occur in many mental representations" (Fodor, 1987, 137). So as well as my thought that John loves Mary, for example, having representations of John, love and Mary as parts, those representational parts may also occur in many – perhaps infinitely many – other thoughts.

Fodor's formulation is more successful than Lewis' as an explanation of systematicity. The fact that the representational parts of the French flag represent parts of what the whole flag represents, for example, is not sufficient for systematicity because the representational parts of the French flag do not occur in any other flags: if there were other flags composed of the same representational parts, the ability to understand the French flag would suffice for the ability to understand those other flags. In general,

transportability explains systematicity because the ability to understand a representation with transportable parts is inseparable from the ability to understand other representations composed of the same parts.

But Fodor's formulation is no more successful than Lewis' at explaining productivity. Take, for example, a language with a vocabulary consisting of an infinite number of names and an infinite number of one-place predicates. The names and predicates are transportable: every name and every predicate occurs in an infinite number of sentences. But it is not possible to understand an infinite number of sentences in the language unless one already understands an infinite number of the names and predicates, because the total number of sentences one understands is only ever equal to the number of names one understands multiplied by the number of predicates one understands, which is finite whenever they are finite.

Moreover, Fodor's formulation still fails to distinguish the explanation provided by the analogy with language from that provided by the analogy with maps. Not only do maps have other maps as parts, but those same parts can reoccur within many – perhaps infinitely many – other maps. The map of Sydney, for example, is part of the map of New South Wales, part of the map of Australia and part of the map of the world. If the location of Sydney were misrepresented, then the map of Sydney might occur as part of a map of Asia, Africa, America or elsewhere. The representational parts of maps, like words, are transportable, so transportability does not distinguish between the explanation offered by the analogy with language and that offered by the analogy with maps.

The quote from Fodor illustrates that the thesis that thoughts have representational parts which may reoccur as parts of other thoughts is a fairly standard formulation of the language of thought hypothesis. I take it to be striking and surprising – even if there turns out to be further relevant differences between maps and language – that this formulation of the hypothesis does not exclude the purportedly alternative hypothesis that the structure of thought is map-like. If the explanation offered by the language of thought

hypothesis is genuinely distinct from that offered by the analogy with maps, then it must be a stronger hypothesis than the hypothesis it is usually characterized as.

V

The productivity of language is explained not merely by the idea that sentences have transportable parts, but also by recursive definitions of what infinite numbers of sentences mean in terms of what finite numbers of their parts mean. The meanings of 'the father of Annette', 'the father of the father of Annette', 'the father of the father of the father of Annette', 'the father of the father of the father of the father of Annette', ... and so on, for example, may be recursively defined by stating that 'Annette' refers to Annette and that 'the father of' prefixed to a referring term refers to the father of the referent of that term (Davidson, 1967, 18).

This suggests that the hypothesis that thought is language-like differs from the hypothesis that it is map-like over whether or not the ability to entertain an infinite number of thoughts should be explained by a recursive definition of what they represent in terms of what a finite number of their parts represent. So the language of thought hypothesis, according to this suggestion, is the hypothesis that the ability to entertain an infinite number of thoughts is explained by a recursive definition of what they represent in terms of what a finite number of their parts represent: I'm able to have an infinite number of thought's about Annette's forefathers, for example, because what is represented by an infinite number of thought's about Annette's forefathers is recursively definable in terms of what is represented by thoughts about Annette and thoughts about fathers.

The hypothesis that there's a recursive definition of what an infinite number of thoughts represent in terms of what a finite number of their parts represent, unlike the hypothesis that thoughts have transportable representational parts, explains productivity. It's not possible to understand an infinite number of sentences in a language consisting only of an infinite number of names and predicates, for example, because what an infinite number of the sentences in that language mean cannot be recursively defined in terms of what a

finite number of their parts mean. Rules stating the meaning of each part – whether recursive or not – will only suffice to specify the meaning of a finite number of sentences, since each part is a component of only a finite number of sentences.

But it's far from obvious that this formulation of the language of thought hypothesis does distinguish it from the hypothesis that thought is map-like. Just as the systematicity of maps may be explained by their transportable parts, the productivity of maps may be aptly explained by recursive definitions of what infinite numbers of maps represent in terms of what finite numbers of their parts represent. The heights represented by contour lines on a topographical map, for example, are defined by base rules stating the height of some of the lines (including the points representing peaks), and a recursive rule stating that the height of each adjacent line is a certain interval more or less (further rules are required to establish which) than the height of its neighbour.

The recursive rule stating that lines represent heights an interval more or less than the heights represented by adjacent lines can define what is represented by an infinite number of lines and thus an infinite number of maps, even when combined with just a finite number of base rules. Combined with a base rule that specifies the height represented by a small circle, for example, the recursive rule defines what height is represented by a circle surrounding that circle, by a circle surrounding the circle surrounding that circle, and so on for any number of circles (with an extra rule to determine whether the heights are ascending or descending). So the heights represented by an infinite number of circles may be defined in terms of the height represented by one.

As well as explaining productivity, recursive definitions of what maps represent in terms of what a finite number of their parts represent explain systematicity compositionally, since the base and recursive rules of the definition state what parts of the map represent, so that anyone who has the abilities corresponding to the base and recursive rules of the definition will have the ability to understand any map composed of those parts. Anyone who has the ability to understand a contour map consisting of seven concentric circles, for example, will have the ability to understand a contour map consisting of just the six

inner circles, because (as long as a base rule states the height represented by one of the six remaining circles) the rules required to define what heights the seven circles represent are sufficient to define what heights the six circles represent.

Before considering objections to the hypothesis that what an infinite number of maps represent should be defined recursively in terms of what a finite number of their parts represent, I want to emphasise an interesting point which is independent of this issue. Proponents of the language of thought hypothesis often claim that only systematicity and not productivity are required to establish it (Fodor, 1987, 148). This is important because productivity, unlike systematicity, is an idealisation: I may be able to think of Annette, Annette's father, Annette's father's father, Annette's father's father's father, and Annette's father's father's father's father, but it's extremely plausible that the length of these thoughts must eventually surpass my ability to entertain them.

If the language of thought hypothesis differed from the analogy with maps by claiming that thoughts, as well as possessing transportable representational parts, are recursively structured, then the argument for the language of thought hypothesis would require the premise that thought is productive, as well as the premise that thought is systematic. The reason for this is that only the hypothesis that thoughts have transportable representational parts is required to explain systematicity; the stronger hypothesis that what an infinite number of thoughts represent should be defined recursively in terms of what a finite number of their parts represent is not required to explain systematicity, but only to explain productivity.

VI

The most important reason for suspecting that maps lack recursive structure is that maps tend to have many more representational parts than sentences. As David Braddon-Mitchell and Frank Jackson write "...there is no preferred way of dividing the map into basic representational units. There are many jigsaw puzzles you might make out of the map, but no one would have a claim to have pieces that are all and only the most basic

units. The reason is that there is no natural *minimum* unit of truth-assessable representation in the case of maps" (Braddon-Mitchell & Jackson, 1996, 180). If maps lack minimum representational parts, their meanings cannot be recursively defined in terms of the meanings of those parts.

The thesis that maps lack minimum representational parts suggests that what maps represent cannot be defined recursively without regress. Although one can, for example, define what the world map represents in terms of what is represented by the maps of the hemispheres, and define what the maps of the hemispheres represent in terms of what is represented by the maps of the demi-hemispheres, and define what the maps of the demi-hemispheres represent in terms of what the maps of the demi-demi-hemispheres represent, this process would continue forever, since the recursive rules stating what maps represent in terms of what their parts represent would never bottom out in base rules stating what their smallest parts represent.[3]

One clarification. The idea that maps lack minimum representational parts, like productivity, is an idealisation: although the world map divides into the maps of the hemispheres, the maps of the hemispheres into the maps of the demi-hemispheres, and the maps of the demi-hemispheres into the maps of the demi-demi-hemispheres, the map will eventually be divided into parts that are either indivisible or too small to be representational. But this does not undermine the underlying point: even if it shows that maps do in fact possess very small minimum representational parts, these minimum representational parts would still be too numerous to explain the ability to understand what is represented by the whole map.

Nevertheless, the point that maps lack minimum representational parts doesn't show that the correct definition of what maps represent is not recursive, because the regress ensues only if the base rules of the definition must state what is represented by a map's

[3] Currie (1995, 130) and Scruton (1983, 107) exploit the same argument to deny, respectively, that cinema and painting are language-like.

minimum representational parts, but this need not be so: it's also possible to define what maps and their parts represent using recursive definitions with base rules that state what their medium sized parts represent. In the recursive definition of the heights represented by contour lines, for example, both the recursive rule stating the interval between lines and the base rules stating the heights of some of the lines state what is represented by medium, rather than minimum, parts of the map.

Two caveats are necessary. First, the lack of minimum representational parts doesn't show that the correct definition of what maps represent has no base and recursive clauses which state what maps represent in terms of what their parts represent, but it does show that those clauses cannot define what is represented by every part of the map, since if the base clauses state what is represented by medium parts of the map directly and the recursive clauses state what is represented by larger parts of the map in terms of what the parts of those larger parts represent, then neither the base nor the recursive clauses will entail what is represented by parts of the map smaller than the parts described by the base clauses. So what smaller parts of the map represent must be defined differently.

Second, the presence of a scale in many maps introduces another explanation of productivity, apart from recursivity. If there were just two unrepeatable icons, a scale could explain the possibility of understanding an infinite number of maps without recursion, because any variation – no matter how slight – in the distance between the two icons would result in a map which represents a different distance obtaining between what the icons represent. This differs from the way recursive definitions explain productivity, because recursive rules have to be reapplied to their own output to produce infinite numbers of cases, whereas single applications of the rule provided by the scale to icons different distances apart may produce infinitely many instances, without recursion.

So the structure of map-like – like language-like – representation is compositional and recursive, but the recursive structure of map-like representation explains only some of the productivity of map-like representation. This suggests that the analogy between maps and thought supports rather than undermines the hypothesis that mental structure is

compositional and recursive. But it also means that it supports it only in a qualified way: the analogy with maps makes it plausible that the systematicity of thought is explained by its compositionality and that some – but not necessarily all – of the productivity of thought is explained by its recursive structure. Thought, according to the analogy, is wholly compositional, but may be only partly recursive.

VII

Braddon-Mitchell and Jackson might agree that recursive definitions of what maps represent in terms of what their medium sized parts represent are possible, but deny – as the quoted passage illustrates – that those definitions are correct, on the grounds that they are arbitrary. A similar problem is raised by Roberto Casati and Achille Varzi, who ask "Suppose you have a uniformly coloured map region: is it composed of its left and right halves or is it composed of its top and bottom halves?" (Casati and Varzi, 1999, 191; see also Fodor, 2007, 108; 2008, 173). The meaning of 'Theaetetus flies', in contrast, should obviously be defined in terms of the meaning of 'Theaetetus' and the meaning of 'flies', so compositional and recursive definitions of what maps represent seem arbitrary in a way that compositional and recursive definitions of what sentences mean do not.

However, just as there are multiple ways of recursively defining what maps represent, there are multiple ways of recursively defining what sentences mean. Quine, for example, has famously pointed out that two alternative recursive definitions may differ such that "….According to one of these systems, the immediate constituents of a certain sentence are 'AB' and 'C'; according to the other system they are 'A' and 'BC'. … which is right?" (Quine, 1970, 392). Quine's example suggests that just as there is no non-arbitrary way of recursively defining what is represented by maps in terms of what their parts represent, there is no non-arbitrary way of recursively defining what sentences mean in terms of what their parts represent either.

Quine draws the conclusion that there is no uniquely correct way to recursively define the meaning of sentences in a language: the definition which gives the meaning of 'John

loves Mary' in terms of the meanings of 'John loves' and 'Mary' and the definition which gives the meaning of 'John loves Mary' in terms of 'John' and 'loves Mary', for example, are equally correct. If one agrees with Quine's conclusion, then the fact that there is no non-arbitrary way to compositionally and recursively define what is represented by a map in terms of what is represented by its parts is of no consequence: all that the explanation of productivity and systematicity requires is that there is at least one such definition, which is not in doubt.

However, there is a non-arbitrary way to choose between alternative recursive specifications of the meaning of sentences in a language: the correct specification is the specification which accurately reflects the tacit knowledge speakers possess of how to understand their language. In particular, each rule in the correct definition of what the sentences of a language mean should correspond to a distinct ability of speakers and interpreters of the language (Evans, 1981, 328). Take, for example, a language consisting of just ten names and ten predicates, which may be combined to form one hundred sentences. The meaning of sentences in this language can be given by a definition with twenty rules – one for each name and predicate – or by a definition with one hundred rules – giving the meaning of each sentence separately (Evans, 1981, 328-330).

Which definition is correct depends on the structure of the language's speakers' and interpreters' dispositions to judge what sentences of the language mean. If the definition with one hundred separate rules is correct, then speakers and interpreters would have one hundred distinct dispositions to judge what each sentence represents. But if the definition with twenty rules is correct then speakers and interpreters of the language would have only twenty distinct dispositions, and the judgement of what any one sentence represents would involve the manifestation of at least two of these, one corresponding to each word. Which dispositions a speaker or interpreter has is constituted by their mind's structure, but indicated by which are lost and acquired independently (Evans, 1981, 330-333).

The same points apply to definitions with recursive rules; the definition is correct if, as well as giving the right meanings for each sentence, each rule corresponds to a distinct

dispositions of speakers to judge what sentences of the language mean (Evans, 1981, 334). English speakers and interpreters, for example, have a single disposition corresponding to the rule that 'the father of' prefixed to a referring term refers to the father of the referent of that term: the disposition to judge for all terms whose reference the speaker or interpreter knows, that 'the father of' prefixed to that term refers to the father of the referent of that term. A speaker who knows that 'Annette' refers to Annette, for example, will be disposed to judge that 'the father of Annette' refers to the father of Annette.

The same reply which Evans makes to Quine on behalf of language can be made on behalf of maps: although there are multiple ways of recursively defining what maps represent in terms of what their parts represent, there is a non-arbitrary way to choose which of these definitions is correct, because the correct definition should reflect the structure of map interpreters' dispositions to judge what maps represent. So the recursive definition of the heights represented by contour lines, for example, is correct because map interpreters have a distinct disposition corresponding to the rule that contour lines represent a height an interval more or less than the height represented by their neighbour: the disposition to judge for any line which the interpreter knows which height it represents, that adjacent lines represent an interval more or less than that height.

VIII

So far I've argued – with two caveats – that the hypothesis that mental structure is language-like and the hypothesis that it's map-like both support the same explanation of the systematicity and productivity of thought: that mental structure is compositional and recursive. But it's important to emphasise that I have not and need not have argued that there's no difference at all between the hypothesis that the structure of thought is map-like and the hypothesis that it is language-like. Rather, I have only argued that – with two caveats – there are no differences between the two hypotheses which are relevant to the explanation of productivity and systematicity, because the productivity and systematicity of both kinds of representation is explained by their compositionality and recursivity.

One way to conclude is not to argue, as I have done, that the analogy with maps supports rather than undermines the language of thought hypothesis, but instead to argue that although the analogy with maps supports the conclusion that the structure of mental representation is compositional, it undermines the grounds for glossing this conclusion as a version of the language of thought hypothesis. Elisabeth Camp, for example, writes, "The premise that I want to challenge is … the claim that any representational system composed of discrete parts with systematic combinatorial rules is a language" (Camp, 2007, 152). Maps, according to Camp, are a counterexample to this premise: although thought is compositional, it does not follow that thought is linguistic.

Given that maps are compositional but non-linguistic, Camp is right that thought's being compositional is no ground for calling thought linguistic. But there is also a more substantive issue; Camp argues that the hypotheses that thought is language-like and that it's map-like offer distinct explanations of systematicity, because maps possess a different kind of compositionality than language. As she writes, "…maps represent by exploiting isomorphisms between the physical properties of vehicle and content …the syntactic principle that combines constituents in maps relies on a fairly direct, albeit selective, isomorphism…." (Camp, 2007, 158). Because maps resemble what they represent, according to Camp, a different kind of compositionality is involved in maps.

I agree that maps differ from sentences due to their resemblance to what they represent, but disagree that these resemblances are involved in a different kind of compositionality. It's helpful to compare the role of resemblance in map-like representation with the role of convention in language-like representation. Even if both thought and language are compositional, the language of thought still differs from ordinary language in that ordinary language is mediated by convention, whereas the language of thought is not. Nevertheless, the fact that the language of thought is not conventional does not show that thoughts possess a different type of compositionality from ordinary language. Nor does the resemblance of maps to what they represent.

To see this, it's worth distinguishing between descriptive and foundational questions about language. Whether or not a language is compositional is a descriptive question about its syntax and semantics; whether or not a language is mediated by convention is a foundational question about what makes it the case that a population speaks that language (Lewis 1969, 204; Stalnaker, 1984, 32-35; 1997, 166-168). The same language with the same syntax and semantics could be conventional for one population, but psychologically necessary for another: in both cases the systematicity and productivity of that language would be explained by its compositional and recursive structure, regardless of what makes it the case that the population speaks that language.

Similarly, whether or not a representational system is mediated by resemblance is a foundational question about what makes that system a population's system of use, not a descriptive question about the system's syntax and semantics. In the system of quotation, for example, words and letters are used to represent themselves. Quotation is a conventional system in our population, but if the resemblance of quotations to themselves were particularly perspicuous, then the same system might have been mediated by resemblance. In either case, whether the system of quotation is mediated by convention, resemblance or psychologically necessity makes no difference to its syntax and semantics (Blumson (2008) defends this point in more detail).

Since the resemblance of maps to what they represent is visually perspicuous – it's this resemblance which makes the system of maps more useful than lists of coordinates and explains our reason for adopting it – the system of map-like representation is mediated by resemblance. But the system of maps would have the same kind of compositional and recursive structure even if we failed to find the resemblance of maps to what they represent perspicuous, but instead used the system of maps in accordance with convention or psychological necessity. Since the question of whether map-like representation is mediated by resemblance is not a descriptive question about the syntax and semantics of maps, it doesn't bear on the kind of structure which maps possess.

Of course, there are obviously some differences between the compositional structure of

maps and sentences: maps are, for example, almost always composed out of spatial parts, whereas sentences may be composed out of spatial or temporal parts. Moreover, the resemblance of maps to what they represent bears on descriptive questions about maps, since it prevents maps from representing anything which they cannot resemble in relevant, usually spatial, respects. Nonetheless, these differences don't seem relevant to the debate over the language of thought hypothesis, since the compositional parts of thoughts are unlikely to be either spatial or temporal and the resemblances between thoughts and what they represent, if thought is mediated by resemblance, are likely to be too abstract to constrain what thoughts may represent.

IX

Another way to distinguish between the map and language of thought hypotheses while conceding that maps have compositional structure is to argue that maps differ from language because maps lack predicative and logical structure (Rescorla, 2009a; 2009b). Michael Rescorla (2009a), for example, exploits Casati and Varzi's (1999, 187-96) compositional theory of (simplified) maps to argue that maps don't feature predication. Rescorla also argues that maps differ from language because they lack logical structure (2009b). In this section I contest the part of Casati and Varzi's theory of maps which supports Rescorla's claim that maps don't feature predication, and use Casati and Varzi's theory to argue that maps, like language, do possess logical structure.

According to Casati and Varzi's theory, what a map represents depends compositionally on what is represented by its atomic map stages. A map stage, according to Casati and Varzi (1999: 192), is any colouring of a map's regions. A map stage is atomic if and only if it colours all and only the regions of a single shade (Casati & Varzi 1999: 192). If, for example, the whole of the region representing France is coloured purple, and nowhere else on the map is purple, then the colouring of that region is an atomic map stage. In contrast, if the region representing Vichy France is coloured purple, and the region representing occupied France is coloured red, then the red and purple colouring of the region representing France is not an atomic map stage.

Colours, according to Casati and Varzi (1999: 191), are like predicates: they represent properties. Regions, according to Casati and Varzi (1999: 191) are like names: they represent objects. Thus, atomic map stages are analogous to atomic sentences. An atomic map stage is true, according to Casati and Varzi's (1999: 194) semantics, if and only if (a) it colours a region of the map which represents a region of the world which has the property represented by its colour and (b) the region of the world represented by the rest of the map does not have that property. So the maximal blue colouring of the map, for example, is true if and only if the region it colours represents a region covered by ocean, and the rest of the map doesn't.

Whole maps are colourings composed of atomic map stages so, according to Casati and Varzi, the truth-conditions of whole maps can be specified recursively in terms of the truth-conditions of their atomic stages. So all other map stages – including whole maps – are true, according to Casati and Varzi (1999: 195), if and only if all their atomic stages are true. The world map, for example, is true if and only if its maximal green and blue colourings are both true. So Casati and Varzi's semantics for maps supports the thesis that maps, like language, have compositional and recursive structure. If Casati and Varzi are right, the analogy with maps supports the same explanation of mental systematicity and productivity as the analogy with language does.

Rescorla agrees that Casati and Varzi's semantics supports the thesis that the structure of maps is compositional and recursive (2009a, 180), but argues that clause (b) in Casati and Varzi's statement of the truth-conditions of atomic map stages undermines the analogy Casati and Varzi draw between colours and predicates. The corresponding statement of the truth-conditions of atomic sentences should, for example, be something like: a predicate concatenated with a name is true if and only if the name refers to a thing which has the property represented by the predicate. 'Snow is white', for example, is true if and only if snow – the referent of 'snow' – has the property of being white – the property represented by 'is white'.

So if colours were truly analogous to predicates and atomic map stages truly analogous to atomic sentences, the correct statement of the truth-conditions of atomic map stages would simply be: an atomic map stage is true if and only if it colours a region of the map which represents a region of the world which has the property represented by its colour (Rescorla, 2009a, 181). So the maximal blue colouring of a map, for example, is true if and only if the region it colours represents a region with the property of being covered by ocean, regardless of whether anywhere else represented by the map has that property. This statement of the truth-conditions of atomic map stages restores the analogy between colours and predicates, and I believe it is correct.

Rescorla agrees that dropping clause (b) from Casati and Varzi's statement of the truth-conditions of atomic map stages restores the analogy between colours and predicates, but argues that the result states the truth-conditions of atomic map stages incorrectly. Imagine, for example, that the property of being mountainous is represented by the colour grey and that the maximal grey colouring of a map of a mountain range has two disconnected parts. Intuitively, the gap between the two disconnected parts of the grey colouring represents a gap in the mountain range. But without clause (b) in the statement of the truth-conditions of atomic map regions, it's compatible with the map that the region represented by the gap could be as mountainous as the rest (Rescorla, 2009a, 187).

Rescorla considers the objection, which I agree with, that it isn't part of the map's truth-conditions that there's a gap in the range, but that this is merely conversationally implicated by the map, since if there were mountains in the region represented by the gap in the grey coloured region, a cooperative map maker would colour that region grey (2009a, 192). Since the map's maker has not coloured that region grey, we're entitled to infer there's a gap in the range. Likewise, if we were considering a journey to Spain or Portugal, and you said 'Portugal is too mountainous', then I'd be entitled to infer that Spain is not too mountainous, even though this is not part of the truth-conditions of what you've said, but merely something conversationally implicated.

But Rescorla argues that conversational implicature can't explain the difference between a map of the mountain range in which grey is absent from some parts and one from which, because the map maker has cut out the non-grey area, those parts of the map are simply missing. Whereas the former map implicates there is gap in the mountain range the latter, since it doesn't represent non-mountainous regions at all, has no such implication. But conversational implicature, according to Rescorla, cannot explain this difference, because if clause (b) is not included in the statement of their truth-conditions, the two maps have exactly the same truth-conditions, and so should, according to Rescorla, generate exactly the same conversational implicatures (Rescorla, 2009a, 193).

But although the two maps have exactly the same truth-conditions, it doesn't follow that they have the same conversational implicatures. Because the former map does represent regions which it does not represent as mountainous, one is entitled to infer that the map maker would have coloured the parts of the map which represents those regions grey if they were mountainous. But since the latter map does not represent those regions at all, one cannot make the same inference in the latter case – the map maker must simply not be interested in those regions, or must not know enough about them to try to represent them. So although the two maps do have the same truth-conditions, the difference in the regions they represent explains the difference between them, without need of clause (b).

Similarly, if we were considering a journey to Spain or Portugal, and I said 'Portugal is too mountainous', then I would not conversationally implicate that Australia is not too mountainous, since Australia is not part of our conversation. But if we were considering a journey to Spain or Portugal, and I said 'Portugal is too mountainous and Australia is Australia', then I might conversationally implicate that Australia is not too mountainous, since – because I'm talking about Australia and in a position to know if it's mountainous – I would mention it otherwise. Since 'Portugal is too mountainous' and 'Portugal is too mountainous and Australia is Australia' have exactly the same truth-conditions, sentences with the same truth-conditions can generate different conversational implicatures.

As well as denying that maps feature predication, Rescorla denies that maps have logical structure (2009b, 392-396). But Casati and Varzi's semantics for maps, which Rescorla relies on in his argument that maps don't feature predication, supports the hypothesis that maps do have logical structure. Casati and Varzi's semantics predicts, for example, that blank maps are tautologies. This is because a map is true, according to Casati and Varzi, if and only if all its atomic regions are true. Since blank maps have no atomic regions – they have no maximally coloured regions since they have no regions of any colour at all – this condition is vacuously met, and blank maps are vacuously true. This prediction is intuitive, since blank maps, like linguistic tautologies, convey no information.

Just as a conjunction is true if and only if all its conjuncts are true, a map is true if and only if all its atomic map stages are true, so whole maps are analogous to conjunctions of their atomic map stages. The world map, for example, is the conjunction of its maximal green and blue colourings. For the same reason, maps entail and are entailed by their atomic map stages: since a map is true only if all its atomic regions are true, the truth of the map entails the truth of the atomic regions, and since a map is true if all its atomic regions are true, the truth of the atomic regions entails the truth of the map. The world map, for example, entails and is entailed by its maximal green and blue colourings. Maps, like language, feature conjunction and entailment, and so possess logical structure.

It might be objected that although maps feature conjunction, they don't feature all logical connectives: although conjunction is present, negation and disjunction, for example, are notably absent. The fact that language is truth-functionally complete – since all classical truth-functional connectives are definable in terms of negation and conjunction – whereas maps are not truth-functionally complete – since not every truth-functional connective is definable in terms of just conjunction – might seem to be a disanalogy between maps and language which not only distinguishes the map from the language of thought hypothesis, but also undermines the ability of the map hypothesis to explain our ability to entertain negative and disjunctive thoughts (Devitt, 2006, 146-7; Rescorla, 2009a, 398).

But this feature of maps is a mere contingency, because other pictorial systems are truth-functionally complete. In Peirce's system of graphs, for example, the conjunction of two statements is represented by their juxtaposition and the negation of a statement by a circle, or cut, around it. Since every classical truth-functional connective is definable in terms of negation and conjunction, Peirce's graphs are truth-functionally complete. A disjunction, for example, consists of each of the disjuncts circled, juxtaposed and then circled again. A material conditional is represented by the negation of the conjunction of the antecedent and the consequent's negation or, in other words, a circle surrounding the juxtaposition of the antecedent with a circle around the consequent (Shin, 2002, 38).

Although not conceding that maps have logical structure, Rescorla also argues that "Even if we were to concede that the representations [maps] in question have logical form, computations defined over them do not exploit it through logical inference" (2009b, 394). In particular, Rescorla argues that robot, and perhaps animal, navigation is not driven by deductive inference in a mental language, but by inductive inference in a map-like format. Robot navigation, according to Rescorla, involves an assignment of probabilities to maps, which the robot updates in accordance with evidence it gathers as it moves around. This explanation, Rescorla suggests, "... does not overtly feature the basic elements of logical form: sentential logical connectives, quantifiers, or even predication" (2009b, 392).

But insofar as Rescorla's suggested explanation of navigation differs from explanations in terms of deductive inference in a language of thought, this is not due to its invocation of maps, but to its appeal to probability. An explanation of robot navigation involving an assignment of probabilities to sentences, propositions or possible worlds, which the robot also updates in accordance with evidence it gathers as it moves around, could proffer the same explanation. Of course, an explanation in terms of an assignment of probabilities to sentences would involve logically structured representations. But if Rescorla's explanation of robot navigation in terms of maps makes no appeal to their logical structure, then this explanation need not appeal to the logical structure of sentences either.

The analogy with maps does not undermine the argument from the best explanation for the hypothesis that thought has compositional and recursive structure because maps, like language, turn out to have compositional and recursive structure as well. Likewise, the analogy with maps cannot undermine arguments for the hypothesis that thought features predication and logical structure, because maps, like language, feature predication and logical structure as well. In general, cartographic and linguistic representation are much more alike than is generally realised, so analogies between maps and thought are much less likely to undermine hypotheses supported by analogies between language and thought than intuition would suggest.

X

The best explanation argument for the language of thought hypothesis is so surprising partly because it purports to establish so much about the structure of thought merely from reflection on the platitudinous facts that thought is systematic and productive. The analogy with maps is supposed to show that although reflection on systematicity and productivity does establish that the mind has some structure or other, very little is knowable from the armchair about that structure: it is possible, for all we know, that mental structure is language-like, but it is equally possible, for all we know, that mental structure is map-like. The structure of the mind, according to this position, is very much an open question, which can only be settled by empirical investigation.

But I have argued that, even if the hypothesis that mental structure is map-like is true, it establishes just as much about the structure of thought as the language of thought hypothesis does. In particular, the analogy with maps still establishes that parts of thoughts represent parts of what whole thoughts represent and that representational parts of thoughts are transportable: they are able to reoccur as parts of infinitely many other thoughts. Moreover, since the productivity of maps is often, though not always, explained by tacit knowledge of a recursive definition of what maps represent in terms of what is represented by a finite number of their parts, the analogy with maps also supports the

hypothesis that mental structure, like linguistic structure, is recursive.

Finally, maps display the same kind of compositional and recursive structure as language. Although what maps represent is mediated by resemblance whereas what sentences mean is mediated by convention, it doesn't follow that the compositionality of maps is different in kind to the compositionality of language, since the resemblance of maps to what they represent is not a descriptive feature of maps, but a foundational feature which explains why they have the descriptive features they do. Likewise, maps and language feature predication and logical structure, so the analogy with maps supports the same hypotheses about the structure of thought as the analogy with language does.

References

Armstrong, David (1973), *Belief, Truth and Knowledge* (Cambridge: Cambridge University Press).

Bermudez, Jose Luis (1983), *Thinking Without Words* (Oxford: Oxford University Press).

Blumson, Ben (2008), "Depiction and Convention", *dialectica* 62(3): 335-348.

Braddon-Mitchell, David and Frank Jackson (1996), *Philosophy of Mind and Cognition* (Cambridge, MA: Blackwell Publishers).

Camp, Elisabeth (2007), "Thinking With Maps", *Philosophical Perspectives* 21(1): 145-182.

Casati, Roberto and Achille Varzi (1999), *Parts and Places: The Structures of Spatial Representation* (Cambridge, Mass.: MIT Press).

Currie, Gregory (1995), *Image and Mind: Film, Philosophy and Cognitive Science* (Cambridge: Cambridge University Press).

Davidson, Donald (1967), "Truth and Meaning", *Synthese* 17: 304-23. Reprinted in Davidson (1984): 17-36.

Davidson, Donald (1984), *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press).

Davies, Martin (1992), "Aunty's Own Argument for the Language of Thought" in J. Ezquerro and J. M. Larrazabal (eds.), *Cognition, Semantics and Philosophy: Proceedings of the First International Colloquium on Cognitive Science* (Dordrecht: Kluwer Academic Publishers).

Devitt, Michael (2006), *Ignorance of Language* (Oxford: Oxford University Press).

Evans, Gareth (1980), *The Varieties of Reference* (Oxford: Oxford University Press).

Evans, Gareth (1981), "Semantic Theory and Tacit Knowledge" in S. Holtzman and C. Leich (eds.) *Wittgenstein: To Follow a Rule* (London: Routledge and Keegan Paul). Reprinted in Evans (1985).

Evans, Gareth (1985), *Collected Papers* (Oxford: Oxford University Press).

Fodor, Jerry (1975), *The Language of Thought* (New York: Crowell).

Fodor, Jerry (1987), *Psychosemantics: the Problem of Meaning in the Philosophy of Mind* (Cambridge, Mass.: MIT Press).

Fodor, Jerry (2007), "The Revenge of the Given", in Brian McLaughlin and Jonathan Cohen (eds.), *Contemporary Debates in Philosophy of Mind* (Oxford: Blackwell Publishers).

Fodor, Jerry (2008), *LOT 2: The Language of Thought Revisited* (Oxford: Oxford University Press).

Lewis, David (1975), "Languages and Language", *Minnesota Studies in the Philosophy of Science* VII: 3-35. Reprinted in Lewis (1983): 163-88.

Lewis, David (1983), *Philosophical Papers I* (Oxford: Oxford University Press).

Lewis, David (1994), "Reduction of Mind", in Samuel Guttenplan (ed.) *A Companion to Philosophy of Mind* (Oxford: Blackwell Publishers). Reprinted in Lewis (1999).

Lewis, David (1999), *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press).

Lycan, William (1993), "A Deductive Argument for the Representational Theory of Thinking", *Mind and Language* 8(3): 404–22.

O'Brien, Gerard and Jon Opie (2004), "Notes Towards a Structuralist Theory of Mental Representation" in H. Clapin, P. Staines and P. Slezak (eds.) *Representation in Mind: New Approaches to Mental Representation* (Elsevier).

Quine, W. V. (1970), "Methodological Reflections on Current Linguistic Theory", *Synthese* 21: 386-98.

Ramsey, Frank (1931), "General Propositions and Causality" in *The Foundations of Mathematics* (New York: Harcourt Brace).

Rescorla, Michael (2009a), "Predication and Cartographic Representation", *Synthese* 169: 175-200.

Rescorla, Michael (2009b), "Cognitive Maps and the Language of Thought", *British Journal of Philosophy of Science* 60: 377-407.

Rey, Georges (1995), "A Not 'Merely Empirical' Argument for a Language of Thought", *Philosophical Perspectives* 9: 201-222.

Scruton, Roger (1983), *The Aesthetic Understanding* (Manchester: Carcanet Press).

Shin, Sun-Joo (2002), *The Iconic Logic of Peirce's Graphs* (Cambridge, Mass.: MIT Press).

Stalnaker, Robert (1984), *Inquiry* (Cambridge, Mass.: MIT Press).

Stalnaker, Robert (1997), "Reference and Necessity" in Crispin Wright and Bob Hale (eds.) *Blackwell Companion to the Philosophy of Language* (Oxford: Basil Blackwell). Reprinted in Stalnaker (2003): 165-187.

Stalnaker, Robert (2003), *Ways a World Might Be: Metaphysical and Anti-Metaphysical Essays* (Oxford: Oxford University Press).