Mental Maps and the Language of Thought

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 the structure of maps, like of language, is also compositional and recursive.

I

Belief, according to many, is a map by which we steer (Ramsey, 1931, 146; Armstrong, 1973; Lewis, 1994, 310-311; Braddon-Mitchell & Jackson, 1996, 177-184). One 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 & 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 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). Productivity and

systematicity show that thought is structured, but not how it is structured.

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 sentences involved in the language of thought hypothesis 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 wish to grant for the purposes of this paper that the analogy with maps can explain everything that the language of thought hypothesis can, including inferential connections between thoughts.

Instead of supporting the language of thought hypothesis by arguing that the analogy with maps cannot explain everything that the language of thought hypothesis can, I will argue that the hypothesis that thought is map-like does not in fact offer an 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. Because the systematicity and productivity of both maps and sentences is explained by their compositional and recursive structure, the explanation suggested by maps is not a genuine alternative to the hypothesis that mental structure is language-like.

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 affect 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 may suggest that thought is mediated by resemblance (O'Brien & Opie, 2004), but they do so without undermining the hypothesis that mental structure is compositional and recursive in the same way as language.

Three clarifications. First, the analogy between maps and thoughts is extremely suggestive (Armstrong, 1973; Ramsey, 1931, 146). Nevertheless, this essay is exclusively concerned with how that analogy bears on the argument from systematicity and productivity 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 it might suggest between mental and linguistic representation, I shall argue that the analogy with maps and the analogy with language both support the same explanation of systematicity and productivity: that the structure of mental representation is compositional and recursive.

Second, I define the language of thought hypothesis as the hypothesis that mental structure is compositional and recursive, because this is the version of the hypothesis which the argument from systematicity and productivity is supposed to support. Stronger definitions, which may be supported by other arguments, are possible. But since it is the argument from systematicity and productivity I wish to defend and proponents of the map hypothesis seek to undermine, the relevant version of the hypothesis is the one which the argument from systematicity and productivity purports to support. Likewise, I will not discuss the specific hypotheses – such as functionalism about mental content – proponents of the map hypothesis have in mind, but merely the analogy with maps itself.

Third, I shall agree with Fodor (1975; 1987) and Braddon-Mitchell and Jackson (1996) in treating the argument from systematicity and productivity as an argument to the best explanation. 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), for which the question of whether there are alternative explanations of productivity and systematicity which are equally as good does not arise. Nevertheless, the analogy with maps still arises as a counterexample to glossing the conclusion of these arguments as a version of the language of thought hypothesis (Camp, 2007, 152), an issue I discuss in section eight.

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 many others (Fodor, 1987, 149). Thought is systematic because the ability to produce some thoughts is connected to the ability to produce many others: the ability to think that John loves Mary, for example, implies the ability to think that Mary loves John (Fodor, 1987, 149-51).[1] Productivity is the ability to produce or understand an infinite number of representations. Thought is productive because the thoughts I have produced are just a small fraction of the infinitely many thoughts I am able to produce (Fodor, 1987, 147-8).

Prima facie, both the analogy with maps and the analogy with language seem 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 many other sentences: 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, 149-51).[2] 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

---

[1] Opinion divides on whether the systematicity of thought is known *a priori* (Evans, 1980, 100-104; Davies, 1992) or *a posteriori* (Fodor, 1978): this is relevant to the question, not raised here, of whether the argument from systematicity is abductive or deductive.

[2] The systematicity of both thought and language is controversial (Johnson, 2004), but is granted by proponents of both the analogy with maps and with language, since if thought is not systematic, then the issue of whether its systematicity is best explained in analogy with maps or languages does not arise.

understand some maps is connected to the ability to understand many 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 with maps and language are equally well placed to explain the productivity and systematicity of thought.

Note that systematicity requires only that the ability to produce or understand some representations is connected to the ability to understand many, but not all, others. 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 and innate paint, because thoughts about green ideas and innate paint do not make sense. Similarly, although one can't understand all maps recomposed of arbitrary parts of maps one can understand, this doesn't show maps aren't as systematic as language, since recompositions of arbitrary parts of maps don't always make sense.

It might also be objected that maps are more systematic than language, on the grounds that people are able to understand every rearrangement of the icons on a map, but unable to understand every rearrangement of words in a sentence. However, it's not true that every rearrangement of icons on a map makes sense: if an arrow represents the direction of a one way street, for example, then that arrow only makes sense if it is placed on a representation of a street. Similarly, it might be objected that maps have a greater degree of productivity than language, because map-like representation is analogue, whereas language-like representation isn't. But this does not matter, since as long as recursion is still the explanation of the productivity of maps, the paper's conclusion is unaffected.

III

According to the language of thought hypothesis, thought and language are alike in possessing compositional syntactic structure. This formulation makes it seem obvious that the analogy with maps provides an alternative explanation of mental productivity and systematicity, because it seems obvious that maps lack syntactic structure of any kind. But if 'syntax' is being used liberally enough to leave open the question of whether thoughts have compositional syntax, it should also be used loosely enough to leave open the question of whether maps have compositional syntax. Settling these questions requires defining compositional syntax in terms of a relationship between representations and their parts which, I will argue, turns out to obtain between maps and their parts.

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. 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. The thought that John loves Mary, for example, has parts which represent John, love, and Mary.

But this characterisation of the hypothesis that thought is language-like does not 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 Earth's southern hemisphere is represented by the map's lower half and the Earth's northern hemisphere is represented 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 productivity and systematicity.

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. It follows, according to Lewis, that the French flag is a language-like representation. 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 characterisation is too weak for two reasons: it neither distinguishes the explanations offered by the two analogies nor explains the phenomena.

(Lewis' claim that parts of thoughts represent parts of what thoughts represent is an unusual formulation of compositionality: it's not obvious, for example, that the symbols for disjunction, negation or conjunction represent any part of what sentences containing them represent. Compositionality is more standardly defined as the thesis that the meanings of whole sentences depends upon the meaning of their parts (Szabo, 2004). But the same points apply to this formulation: what whole maps represent is dependent on what their parts represent and the French flag's representing liberty, equality and fraternity depends on its having a part which represents each of them.)

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 representational parts.

But Fodor's formulation is no more successful than Lewis' at explaining productivity. Take, for example, a language with a vocabulary consisting of just 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 language of thought hypothesis from the hypothesis that thought is map-like. 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. Since the representational parts of maps, like words, are transportable, the idea that thoughts have transportable parts does not distinguish between the explanations offered by the analogy with language and 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 distinctions between cartography 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 rules which define what infinite numbers of sentences mean in terms of what finite numbers of their parts mean.[3] 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', 'the father of the father of the father of the father of the father of Annette', … and so on, for example, is 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).

A recursive rule is a rule which applies to its own output; the possibility of defining an infinite number of cases arises from the repeated reapplication of the rule to previous instances. A recursive definition is a definition consisting of a base rule, which specifies initial cases, and a recursive rule, which specifies subsequent cases: The natural numbers, for example, can be defined recursively by a base rule which states that zero is a natural number and a recursive rule which says that the successor of a natural number is a natural

[3] It's debated whether recursion is essential and unique to human language (Hauser, Chomsky & Fitch, 2002; Pinker & Jackendoff, 2005); this raises issues which are largely orthogonal to the debate over the language of thought hypothesis.

number: one is natural because one is the successor of zero, two is natural because two is the successor of one, three is natural because three is the successor of two, … and so on.

This suggests that the hypothesis that thought is language-like may differ from the hypothesis that thought is map-like over whether or not what thoughts represent should be defined recursively in terms of what a finite number of their parts represent. The language of thought hypothesis, according to this suggestion, is the hypothesis that the correct definition of what thoughts represent is a recursive definition in terms of what a finite number of their parts represent: I am 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 defined recursively in terms of what is represented by thoughts about Annette and thoughts about fathers.

The hypothesis that there is a correct 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 the 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 only a finite number of sentences is composed of each part.

Two clarifications. First, it's probable that there will be multiple recursive definitions of what thoughts, maps or sentences represent in terms of what a finite number of their parts represent. Nevertheless, I will argue below that not all these definitions are correct. So the language of thought hypothesis is the hypothesis that the correct definition of what thoughts represent is a recursive definition in terms of what a finite number of their parts represent. Second, since the productivity of language is explained by recursively defining what an infinite number of sentences mean in terms of what a finite numbers of their parts mean, I take it that this formulation of the language of thought hypothesis is right,

even if it fails to distinguish between the map and language of thought hypotheses.

Unfortunately, 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, is correctly 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 succeeding 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 the systematicity of maps 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 height

the seven circles represent are sufficient to define what height the six circles represent.

Before considering objections to the hypothesis that what an infinite number of maps represent is correctly 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 differs 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 does 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 were 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 a finite number of minimum representational parts, then 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 is represented by the demi-demi-hemispheres, 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.[4]

Two clarifications. First, 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, eventually the map will 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.

Second, it seems that some maps, like those containing icons, do have minimum representational parts. Moreover, it is very difficult to generalise about the semantics of all maps: some maps contain icons, whereas others do not; some maps are drawn to scale, whereas others are not; and the conventions of political maps are very different from those of physical maps. Nevertheless, to illustrate the potential regress, Braddon-Mitchell

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

and Jackson merely require that some maps lack minimum representational parts. So rather than disputing whether Braddon-Mitchell and Jackson are right about all maps, I will grant that some maps lack minimum representational parts, but argue that even what those maps represent can be defined recursively.

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 it is assumed that the base rules of the definition must state what is represented by a map's minimum representational parts, but this need not be so: it's also possible to define what maps and their parts represent using a recursive definition 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.

Recursive definitions of what maps represent in terms of what their medium sized parts represent is consistent with maps lacking minimum representational parts, because such definitions may include recursive rules which define what smaller as well as larger parts represent, and the repeated reapplication of those rules may define what is represented by smaller and smaller representational parts, without cease. The approximate shape represented by a line, for example, can be defined using a base rule directing a new line to be drawn between the furthest points of the original line, and a recursive rule directing that new lines be replaced by lines between their endpoints and the furthest points on the original line (this rule is recursive because it applies to its own output; it directs new lines to be replaced by new lines).

The shape represented by a line to a degree of approximation, according to this definition, is the shape reached by applying these rules until the distances between new lines and the

furthest points on the original line is below that degree of approximation.[5] If the degree of approximation is high, then this definition would deliver that France is hexagonal and Italy is boot-shaped, whereas if the degree of approximation is low, the definition would deliver much more complicated shapes for both countries. Although recursive, the definition accommodates the point that maps lack minimum representational parts, because the definition may be applied using any degree of approximation, no matter how small, so even small variations in the shape of a line affect the shape it represents.

One caveat is required, because the presence of a scale in many maps introduces another explanation of productivity, apart from recursivity. If there were just two unrepeatable icons, the 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 language of thought hypothesis – construed as the hypothesis that mental structure is compositional and recursive – but also that it supports it 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 only mostly recursive.

[5] This is an algorithm used in computerised map production. See Douglas and Peucker (1973), Hershberger and Snoeyink (1992), and Visvalingam and Whyatt (1991).

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 ways that compositional and recursive definitions of the meanings of sentences are 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 of these definitions is correct depends on the structure of the speakers' dispositions to judge what the 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 that for all terms whose reference the speaker or interpreter knows, '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 that for any line which the interpreter knows which height it represents, adjacent lines represent an interval more or less than that height.

VIII

So far I’ve argued – with one caveat – 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 one caveat – 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 two kinds of question about languages. Whether or not a language is compositional is a descriptive question about the syntax and semantics of that language; whether or not a language is mediated by convention is a foundational question about what makes that language a population's language of use (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, which remains the same regardless of whether the language is conventional or psychologically necessary.

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 being 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.

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 particularly useful, 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 likely to be too abstract to constrain what thoughts may represent.

IX

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, but not always, explained by tacit knowledge of a correct recursive definition of what maps represent in terms of what is represented by a finite number of their parts, the analogy with maps is also able to establish that mental structure, like linguistic structure, is recursive.

References

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

Braddon-Mitchell, David and Frank Jackson (1996), *Philosophy of Mind and Cognitive Science* (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).

Douglas, David and Thomas Peucker (1973), "Algorithms for the Reduction of the Number of Points Required to Represent a Digitized Line or its Caricature", *The Canadian Cartographer* 10(2): 112-122.

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).

Hauser, Marc, Noam Chomsky and Tecumseh Fitch (2002), "The Faculty of Language: What Is It, Who Has It and How Did It Evolve?", *Science* 298: 1569-1579.

Heck, Richard (2007), "Are There Different Kinds of Content?" in Brian McLaughlin and Jonathan Cohen (eds.), *Contemporary Debates in Philosophy of Mind* (Oxford: Blackwell Publishers).

Hershberger, John and Jack Snoeyink (1992), "Speeding Up the Douglas-Peucker Algorithm", *Proceedings of the International Symposium of Data Handling* 5: 134-143.

Johnson, Kent (2004), "On the Systematicity of Language and Thought", *Journal of Philosophy* 101(3): 111-39.

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).

Pinker, Steven and Ray Jackendoff (2005), "The Faculty of Language: What's Special About It?", *Cognition* 95: 201-236.

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).

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

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).

Szabo, Zoltan Gendler (2004), "Compositionality", *Stanford Encyclopedia of Philosophy*, URL=<http://plato.stanford.edu/entries/compositionality/>.

Visvalingam, Mahes and Whyatt, Duncan (1990), "The Douglas-Peucker Algorithm for Line Simplification: Re-evaluation through Visualization", *Computer Graphics Forum* 9: 213-228.