Philosophical Perspectives, 21, Philosophy of Mind, 2007

THINKING WITH MAPS*

Elisabeth Camp
University of Pennsylvania

Most of us create and use a panoply of non-sentential representations throughout our ordinary lives: we regularly use maps to navigate, charts to keep track of complex patterns of data, and diagrams to visualize logical and causal relations among states of affairs. But philosophers typically pay little attention to such representations, focusing almost exclusively on language instead. In particular, when theorizing about the mind, many philosophers assume that there is a very tight mapping between language and thought. Some analyze utterances as the outer vocalizations of inner thoughts (e.g. Grice 1957, Devitt 2005), while others treat thought as a form of inner speech (e.g. Sellars 1956/1997, Carruthers 2002). But even philosophers who take no stand on the relative priority of language and thought still tend to individuate mental states in terms of the sentences we use to ascribe them. Indeed, Dummett (1993) claims that it is constitutive of analytic philosophy that it approaches the mind by way of language.

In many ways, this linguistic model is salutary. Our thoughts are often intimately intertwined with their linguistic expression, and public language does provide a comparatively tractable proxy for, and a window into, the messier realm of thought. However, an exclusive focus on thought as it is expressed in language threatens to leave other sorts of thought unexplained, or even to blind us to their possibility. In particular, many cognitive ethologists and psychologists find it useful to talk about humans, chimpanzees, birds, rats, and even bees as employing cognitive maps. We need to make sense of this way of talking about minds as well as more familiar sentential descriptions.

In what follows, I investigate the theoretical and practical possibility of non-sentential thought. Ultimately, I am most interested in the contours of distinctively human thought: what forms does human thought take, and how do those different forms interact? How does human thought compare with that of other animals? In this essay, however, I focus on a narrower and more basic theoretical question: could thought occur in maps? Many philosophers are convinced that in some important sense, thought *per se* must be language-like:

that there are constitutive features of thought which can only be explained if we assume that it has a sentential form. I will argue that on the contrary, these features can also be satisfied in a cartographic representational system. There are good reasons to believe that much of our own thinking is sentential, but these reasons depend on *what* we think—on the particular sorts of contents that we represent and reason about—rather than on general features of thought as such.

§1: Why Must Thought Be Language-Like?

The classic reason for thinking that thought must be language-like is that only this assumption can explain or justify the *systematicity* of thought. The argument has been articulated most prominently by Jerry Fodor (e.g. Fodor 1975, 1987; Fodor and Pylyshyn 1988; Fodor and McLaughlin 1990); but it has also been defended by philosophers of a more rationalist, neo-Kantian orientation. In brief, the argument goes as follows:

1. There are systematic relations among the contents that a thinker can represent and reason about.
2. Systematic relations in content must be reflected by correlative structure in a thinker's representational and reasoning abilities.
3. Structured representational abilities require a system of representational vehicles which are composed of recurring discrete parts combined according to systematic rules.
4. Any system of representational vehicles composed of recurring discrete parts combined according to systematic rules is a language.

Therefore: there must be a language of thought.

I'll consider the premises in turn.

Premise 1

Fodor takes the first premise as an empirical observation: when we examine the minds around us, we find systematic relations among the contents that they can represent. Thus, Fodor & Pylyshyn (1988, 39) write:

> What does it mean to say that thought is systematic? Well, just as you don't find people who can understand the sentence 'John loves the girl' but not the sentence 'the girl loves John', so too you don't find people who can *think the thought* that John loves the girl but can't think the thought that the girl loves John.

Gareth Evans (1982, 100) makes the same basic point in more *a priori* terms, about thoughts themselves:

> It seems to me that there must be a sense in which thoughts are structured. The thought that John is happy has something in common with the thought that Harry is happy, and something in common with the thought that John is sad.

The two states of affairs in the world, of John loving the girl and the girl loving John, are systematically related in the sense that they involve the same individuals and properties: John, the girl, and the relation of loving; the only difference (albeit an important one) is who loves who. Premise 1 amounts to the claim that the abilities to think about such states of affairs cluster together: if you can think one thought about John, Harry, the girl, loving, or being happy, you can also think other thoughts about them, such as that *Harry loves the girl* or that *the girl is happy*.[1] Conversely, there are also systematic limitations in the contents a thinker can represent: a thinker who can't entertain the thought that *John blighted the girl* also can't think that *Harry blighted the girl*, that *the girl blighted the cow*, or anything else about blighting.

These systematic patterns among the contents a thinker can and can't represent are also manifested in how thinkers reason about those contents. As Fodor and Pylyshyn (1988, 46) say,

> It's a 'logical' principle that conjunctions entail their constituents... Correspondingly, it's a psychological law that thoughts that *P&Q* tend to cause thoughts that *P* and thoughts that *Q*, all else being equal.

This point can be extended beyond deductive inference. Thus, for any inductive law—say, that if all heretofore observed *Fs* are *G*, then probably all *Fs* are *G*—there is an analogous psychological law—say, that many thoughts of the form $This_1$ *F is G*, $This_2$ *F is G*..., plus the thought that *I haven't seen any F's that aren't G*, tend to cause the thought that *All Fs are G*. In both cases, transitions between thoughts track relations among represented contents.[2]

In the more rationalist tradition, Tim Crane (1992, 146–7) claims that justifying the transitions thinkers make between thoughts requires us to recognize systematic relations among their contents:

> If we simply wanted to represent facts, then our beliefs would only need to have 'whole' contents. All that would matter would be whether a content was true or false. The *fact* might have constituents (particulars and properties) but they would have no reflection in the content, since (to echo Frege) they would as it were have no role, no 'meaning of their own'. But once we consider the role our beliefs play in reasoning, then it starts to become clear why their contents need constituents. A thinker who believes that *a is F*, and that *b is F*, and that *a is not b* will be disposed to believe that *at least two things are F*. Surely the states in this inference cannot just have unstructured contents, or we would not be able to explain its validity.[3]

Here, Crane, like Evans, talks about "structure" instead of "systematicity," but the basic point again concerns systematic relations among contents: the reason that the transition from believing that *a is F*, *b is F*, and *a is not b* to believing that *at least two things are F* is justified is that there are systematic, truth- and justification-preserving relations in the contents of those beliefs.

Premise 2

Premise 1 claims that there are systematic empirical and/or normative relations among a thinker's abilities to think and reason with whole thoughts, in virtue of what those thoughts are about. Premise 2 claims that in order to explain this systematicity in contents, we must assume that the thinker's representational abilities are themselves structured, in the sense that they must be produced by interacting constituent abilities to represent various parts of the world. Thus, the reason that the ability to think that *John is happy* clusters together with the ability to think that *Harry is happy* is that both thoughts involve exercising a general ability to think about *being happy*. Likewise, the reason that a thinker who makes the transition from the thought that *John is happy* to *John is not sad* also makes the transition from the thought that *The girl is happy* to *The girl is not sad* is that she has a general ability to infer that if someone is happy, then they are not sad.

Fodor and Pylyshyn support this claim by appeal to explanatory parsimony. Unless we posit distinct interacting abilities to represent the objects, properties, and relations that together constitute whole states of affairs, the systematic patterns of abilities and limitations that we observe among the whole contents that a thinker can represent will remain unexplained. More importantly, as the range of contents a thinker can represent increases, it becomes exponentially more efficient to posit distinct, interacting abilities to represent parts of contents rather than distinct, unstructured abilities to represent entire contents. For instance, baboons clearly demonstrate an awareness of all the dominance relations in their troops of approximately 40 animals. In principle, we could explain this by postulating that they have memorized each of the approximately 800 dominance relations separately. But it is much more parsimonious to assume that they represent each of their approximately 40 troopmates, plus a general dominance relation (cf. Cheney and Seyfarth 2007).[4]

Evans argues for structured representational abilities, not on empirical scientific grounds, but by claiming that our ordinary practices of mental-state-ascription commit us to there being a "common explanation" for a thinker's ability to think about related contents (1982, 102). Ascribing the thoughts that John is happy and that Harry is happy, he thinks, both involve ascribing the ability to represent something as *being happy*. But insofar as a thinker really has the ability to think about *being happy,* she should be able to apply this ability in thinking about other individuals as well. Similarly, if she really has the ability to think about John, then she should be able to think of him not just as being happy, but also as being sad, or fat, or bald. This delivers the "Generality Constraint":

> If a subject can be credited with the thought that *a* is *F* then he must have the conceptual resources for entertaining the thought that *a* is *G*, for every property of being *G* of which he has a conception (1982, 104).

Analogously, if a thinker can be credited with making an inference from *John is happy*, *Harry is happy*, and *John is not Harry* to the conclusion that *At least two* *people are happy,* then that thinker must have a general disposition to move from *a is F*, *b is F*, and *a is not b* to *At least two things are F,* for every concept *a*, *b*, and *F* in her possession.

Premise 3

The third premise in the argument for a Language of Thought claims that structured representational abilities must be underwritten by structured representational vehicles: by mental representations which are composed of recurring, systematically interacting parts. As Fodor & Pylyshyn (1988, 39) say, continuing the quote from above:

> But now if the ability to think that John loves the girl is intrinsically connected to the ability to think that the girl loves John, that fact will somehow have to be explained.... For Representationalists... entertaining thoughts requires being in representational states (i.e., it requires tokening mental representations)... [T]he systematicity of thought shows that there must be structural relations between the mental representation[s] that correspond [to the two thoughts]... the two mental representations, like the two sentences, *must be made of the same parts*.

For Fodor and Pylyshyn, the claim that mental representations of related contents "must be made of the same parts" amounts to the claim that at the "cognitive level"—the level of description which specifies how brain states represent information about the world—(a) there must be physical properties which encode each object, property, and relation that enters into those contents; (b) the physical structures among those properties must encode the structural relations among those represented constituents; and (c) these physical structures must cause the overall representational system to behave as it does. Thus, representing that *John is happy* requires that, at the cognitive level of description, there be a physical structure in the brain which combines two distinct physical properties, with the functions of representing *John* and *being happy*, respectively, into a larger structure which encodes the relation of predication. And reasoning from *John is happy* to *John is not sad* must consist in a physical process transforming this physical structure into another one which also involves the physical property that encodes *John,* but now combining it with physical properties that represent *being sad* and *not*.

According to Fodor and Pylyshyn, this conclusion follows from the general scientific principle that sameness and difference of observed effects—here, abilities to represent and reason about objects, properties and relations—entails sameness and difference of unobserved causes—here, physical brain states. However, the cognitive level of description may be quite high-level. In the case of public languages, we classify many different vocalized and written tokens as instances of the same sentence, in virtue of their shared functional properties.

Likewise, the claim that all of a thinker's thoughts about *being happy* are underwritten by a common physical property doesn't entail that a specific set of neurons always and only fires when the thinker thinks *happy* thoughts.

Philosophers in the rationalist tradition tend to be more interested in normativity than in brain states. But many of them also endorse the claim that believing and desiring involve tokening mental representations with recurrent parts, in some sense of those terms; they just prefer to remain neutral about what exactly is involved in 'tokening a mental representation' and 'having parts'. Martin Davies (1991), straddling both traditions, harnesses realism about mental states in support of an *a priori* argument for Premise 3. In order to construe the Generality Constraint in a "full-blooded" way—as requiring that there be a "common explanation" of a thinker's ability to think various thoughts about, say, *being happy*—we must postulate a common cause which underwrites that ability every time it is exercised. But surely, Davies claims, any real cause must ultimately be a physical mechanism. Therefore, he concludes, each of the thinker's conceptual abilities must be underwritten by a distinctive physical brain structure.

Premise 4

The final premise, which is often conflated with the previous one, is that any system which combines recurrent parts according to systematic rules to generate whole representations is a language. Because we are considering a representational vehicle *of thought*, this delivers the conclusion that there is a language of thought.[5] As Fodor and Pylyshyn (1988, 39) say, immediately following the quote above,

> But if this explanation is right (and there don't seem to be any others on offer), then mental representations have internal structure and there is a language of thought.

Likewise, here's Fodor (1987):

> What's at issue ... is the internal structure of these functionally individuated states. Aunty [i.e. philosophical orthodoxy] thinks they have none; only the intentional objects of mental states are complex. I think they constitute a language; roughly, the syntactic structure of mental states mirrors the semantic relations among their intentional objects.

Philosophers in the rationalist tradition also embrace the need for a language of thought. For instance, José Luis Bermudez (2003, 111) starts with something like Crane's claim above—that in order to justify inferences between thoughts we must appeal to systematic relations in their contents—and concludes that genuine reasoning requires a linguistic vehicle:

> We understand inference in formal terms—in terms of rules that operate on representations in virtue of their structure. But we have no theory at all of formal

> inferential transitions between thoughts that do not have linguistic vehicles.... Clearly, it is a necessary condition on there being formal inferential transitions between contentful thoughts that those thoughts should have structured contents. Nonetheless, it is not a sufficient condition. Formal rules of inference do not operate on thought-contents but rather on the vehicles of those contents. They are syntactic rather than semantic.

Although Bermudez doesn't say much to explicitly defend the claim that "we understand inference in formal terms," I believe he's thinking that if we want to explain why, say, a thinker's inference from *John is happy*, *Harry is happy* and *John is not Harry* to *At least two people are happy* is an instance of valid reasoning, then it's not enough to point out that if the first three states of affairs obtain then the last one will also obtain. In addition, this validity must be demonstrable from the thinker's own perspective, given her way of representing those contents.[6] And this requires that the vehicles by means of which the thinker represents those contents must have a form which makes it possible to justify the transition.[7] But because he assumes that only a linguistic system can have formal rules of valid inference, he concludes that justified inference requires a language of thought.[8]

Evaluating the Argument for a Language of Thought

Given the strength and breadth of the argument's conclusion, it's no surprise that each step in the argument has been hotly contested. I'll very briefly review some possible objections to Premises 1 through 3. First, one might dispute the claim that thinkers' representational and reasoning abilities are so systematic.[9] Investigation into animal cognition is often confounded by apparent failures of systematicity: an animal seems to have all the constituent representations it needs to arrive at some further representation which it should be highly motivated to act upon, but it fails to act in the relevant way. Humans also regularly exhibit significant failures of systematicity. For instance, performance on the Wason Card Selection Task, in which people are asked which of four cards they need to turn over to test a material conditional, varies dramatically depending on the subject of the conditional being tested: if the rule concerns social behavior, people perform well, while if it concerns abstract relations between numbers and colors, they perform abysmally (Wason and Johnson-Laird 1972, J. St. B. T. Evans 1982). This looks like a case where both inferences should be "of the same form," and so it seems that according to Fodor and Pylyshyn, the same "psychological law" should apply in both cases.[10]

Second, one might deny that thought must be structured: perhaps a creature could have simple unstructured thoughts, like *Threat!* or *Food!* Third, one might accept that thinkers' representational and reasoning abilities are structured, but deny that there must be a correlatively structured vehicle of thought. This is Evans's position; as he says,

> I do not wish to be committed to the idea that having thoughts involves the subject's using, manipulating, or apprehending *symbols*—which would be entities with non-semantic as well as semantic properties... I should prefer to explain the sense in which thoughts are structured, not in terms of their being composed of several distinct *elements*, but in terms of their being a complex of the exercise of several distinct conceptual *abilities* (1982, 100–101).

We obviously need some account of the causal underpinnings of mental states, abilities, and processes; but many people believe that structured, stable patterns of representation and reasoning could emerge without precisely correlative underlying physical constituents and mechanisms. Likewise, one might deny Bermudez's claim that we must understand inference and reasoning in terms of *formal* relations among representational vehicles. Perhaps it suffices to appeal to substantive relations among represented contents, where "contents" are understood either as possible worlds, or else as structured Fregean or Russellian propositions.[11] Indeed, Gil Harman (1986, 20) has argued that "there is no clearly significant way in which logic is specially relevant to reasoning."

That said, Fodorian computationalism does provide a comprehensible, straightforward model for *a* way the mind/brain might work, which many philosophers and cognitive scientists have found enormously fruitful. It's also important to see what endorsing a structured vehicle of thought doesn't entail. First, it doesn't require that the thinker consciously attends to that vehicle; it is enough for her to represent *with* the vehicle—for it to play the right functional role in her thinking.[12] Nor does granting that a structured vehicle plays an important causal and explanatory role in thought entail that it does all of that work by itself: the vehicle's functional role within the overall cognitive system is equally important. As Pylyshyn says, "The appropriate subject of our analysis of representation should be *not* the representation per se but a *representational system* consisting of the pair (representation, process)" (cited in Anderson 1978, 250).

Weak and Strong LOT

The premise that I want to challenge is the one that draws the least explicit attention: Premise 4, the claim that any representational system composed of discrete parts with systematic combinatorial rules is a language. At most, the arguments I've rehearsed only take us as far as Premise 3. Thus, at best they only establish what we might call *Weak-LOT*: the claim that thought requires a system of representational vehicles with *some* recurrent constituents that can be recombined according to *some* set of rules to produce representations of systematically related entire contents.[13] This falls significantly short of the proffered conclusion, which we might call *Strong-LOT*: that claim thought requires a specifically sentential structure and semantics.

Although Premise 4 is rarely articulated or defended explicitly, people regularly construe "the Language of Thought Hypothesis" as supporting the stronger claim. Bermudez appeals specifically to linguistic structure in the passage about the formal justification of reasoning cited above.[14] Michael Devitt (2005, 146) characterizes the LOT hypothesis as requiring "that the simplest meaningful parts of the representation involved in a thought be like words, and that the structure of the representation be like the syntactic structure of a sentence." And like Bermudez, Devitt (2005, 147) invokes the argument from reasoning in support of a specifically linguistic vehicle:

> Formal logic gives us a very good idea of how thinking might proceed if thoughts are represented linguistically... We still have very little idea how thinking could proceed if thoughts were not language-like but, say, map-like.

Finally, Dummett (1989, 197) claims that "a fully explicit verbal expression is the only vehicle whose structure must reflect the structure of the thought," thereby implicitly assuming that thought itself has a language-like structure.[15]

The assumption that thought is language-like might not seem so contentious if we're only considering human thought: after all, normal humans do often express their thoughts verbally, and often experience the phenomenology of thinking in language. However, the theorists we've considered take themselves to be investigating the structure of thought in general. Fodor and his co-authors repeatedly emphasize that their empirical claims about systematicity extend to non-human animals:

> Linguistic capacity is a paradigm of systematic cognition, but it's wildly unlikely that it's the only example. On the contrary, there's every reason to believe that systematicity is a thoroughly pervasive feature of human and infrahuman mentation (Fodor and Pylyshyn 1988, 37).
>
> It may be partly a matter of taste whether you take it that the minds of animals are productive; but it's about as empirical as anything can be whether they are systematic. And – by and large – they are. (Fodor 1987).

Dummett, Crane, and Bermudez are driven by less empirical considerations: they want to identify a condition on genuine thought, or at least on the sort of conceptual thought that's involved in genuine reasoning. But their identified conditions rely on quite general features of representation and reasoning, and are intended to have commensurately general application.

However, as I noted at the outset, we're all quite familiar with representational systems that appear, at least intuitively, to employ very different combinatorial structures than language. Diagrammatic representational systems, such as Venn diagrams, are formed by combining formal elements like circles, dots, and lines according to systematic rules which determine the representational content of the whole. Further, they are governed by formal rules of inference which are sound and complete, up to expressive equivalence with monadic first-order predicate logic (Shin 1994). But the elements and combinatorial rules for diagrams are very different than those for sentences. Thus, even if we grant the requirement of formal validity, it's simply false that that "we have no theory at all of formal inferential transitions between thoughts that do not have linguistic vehicles," as Bermudez, Devitt, and others claim.

In this paper I want to focus more narrowly on another alternative: maps. Cartographic systems range along a continuum from the nearly pictorial, such as Google's map-satellite hybrids, to the nearly diagrammatic, such as subway maps. In the next section, I show that sentential and cartographic systems do indeed employ different combinatorial principles. Here, I just want to establish that familiar maps, such as Rand-McNally city and road maps, meet the demands that Weak-LOT claims a vehicle of thought must satisfy.

Such maps are clearly constructed out of recurrent formal elements that make a common semantic contribution each time they occur: for instance, on many maps any solid line of a certain width signifies a street, any blue line or blob signifies a river or lake, and any cross signifies a church. Further, the representational import of the entire map is a systematic function of the way in which those elements are combined: if two lines intersect, with a blob in one quadrant and a cross in the other (Figure 1a), then this represents two intersecting streets with a church across from a pond. By contrast, if the two lines are drawn in parallel, with the cross above the blob (Figure 1b), then these same elements represent a different but related situation, in which a church is north of a pond and between two parallel roads.

Because maps' constituents are systematically recombinable, in this way, they also satisfy the Generality Constraint: a cartographic system that enables a thinker to represent the locations of City Hall, the Delaware River, Dunkin' Donuts locations, and bus routes thereby has the representational resources to represent those same features in any spatial configuration. I don't believe that anyone has developed formal rules for reasoning with maps in the way that Shin (1994) has for extended Venn diagrams. And in §3 I'll argue that rules would look very different than those for either sentential or diagrammatic logics. However, I see no theoretical reason why one couldn't define formal updating rules for dynamic reasoning with maps that would mirror semantic changes in the relations among the represented states of affairs, and thus would be reliably and demonstrably truth-preserving. And I believe that such rules could be used in genuine reasoning. Thus, I see no principled reason why maps fail to satisfy the arguments offered above for a compositional system of representational vehicles.

§2: The Syntactic and Semantic Principles of Maps and Sentences

Perhaps the most natural response to my drawing the distinction between Weak-LOT and Strong-LOT is to deny that the distinction is theoretically interesting. All that really matters for a theory of mind, one might argue, is Weak-LOT: that there be some discrete symbols combined according to some set

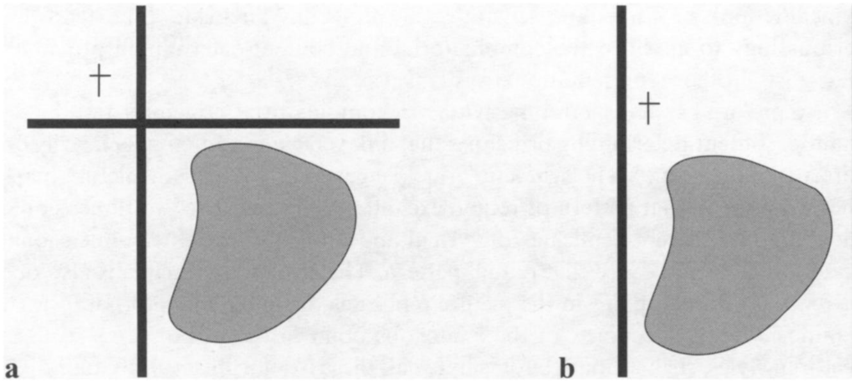

Figure 1. Two maps constructed of the same parts in different ways representing systematically related but distinct states of affairs.

of rules, such that the content of the entire representation is a function of the meaning and mode of combination of those symbols. The specific symbolic and combinatorial principles employed by a representational system are, one might think, at best a topic for merely empirical, neurophysiological investigation. At the extreme, one might insist that diagrams and maps just *are* sentences written in a funny notation. Thus, Eliot Sober (1976, 141) claims that the fundamental distinction between pictures and sentences is that genuine pictures are analog, in the sense that they represent continuous values (e.g. color) in a continuous way; given this assumption, he then claims that "where [picture-like representational systems] are digital, they simply *are* linguistic systems of a certain kind." Likewise, Bermudez (2003, 155) claims that "the essence of language is the combination of symbols with each other to express thoughts, taking thoughts to be complex entities that can be assessed for truth or falsity." Because maps satisfy these conditions, perhaps they don't constitute a counterexample to the claim that thought must be language-like after all.

In this section I argue that maps and languages do operate according to importantly different combinatorial principles, and that as a result, thinking in maps is substantively different from thinking in sentences. In principle, we can distinguish two aspects of any representational system. On the one hand, there is the *form* of its representational vehicles: what the basic representational constituents are and the principles that govern how those constituents are put together. On the other hand, there is their *content*: what those constituents are about and the principles that determine that they are about this. The overall content of a complete representational vehicle is a function of the content of its basic representational parts and the significance of their mode of combination. In linguistic systems, this distinction is clear, and corresponds to the distinction between syntax and semantics. In other systems, the distinction is less clear, because the two principles interact in interesting ways. Although we typically think of syntax and semantics as specifically linguistic, I'll extend this terminology to apply to the combinatorial and content-determining principles that govern other representational systems as well.

In pictorial systems, both the syntactic, combinatorial principles and the semantic, content-determining principles that link vehicle to content rely heavily on direct isomorphism.[16] The syntactic principle generating a realistic picture maps the two-dimensional pattern of retinal excitation onto another two-dimensional medium, and thereby replicates the visual appearance of the three-dimensional scene which would cause that retinal pattern. The semantic principle is also one of replication: each point in the picture replicates the apparent color, or at least luminosity and reflectancy, of the analogous point in the world. (Less realistic pictorial styles, such as impressionism, tweak these isomorphisms.) Because their syntactic and semantic principles both rely on fairly direct replication of the visual appearance of a scene, in pictures the distinction between syntax and semantics is blurry: insofar as we can discern syntactic 'parts' to a picture at all, these are either just points in a two-dimensional array, or else regions whose boundaries are given by salient boundaries in the scene being represented; and in either case, the semantic principle simply replicates the visual appearance of that very point or region.

Because pictorial systems replicate the visual appearance of a scene by largely replicating that visual appearance itself, pictures can only explicitly represent features that are themselves visually perceptible.[17] This also makes them highly analog modes of representation, in two respects: they deliver information about a continuous spatial array, and the information they deliver about each point in that array is itself typically continuous—for instance, specifying fine-grained color (or at least greyscale) values. On the one hand, this rich multi-dimensional specificity enables pictures to communicate lots of information simultaneously in a compact, comprehensible form. Fred Dretske (1981, 137) illustrates the point with a cup of coffee:

> If I simply *tell* you, 'The cup has coffee in it,' this (acoustic) signal carries the information that the cup has coffee in it in digital form. No more specific information is supplied about the cup (or the coffee) than that there is some coffee in the cup. You are not told *how much* coffee there is in the cup, how large the cup is, *how dark* the coffee is.... If, on the other hand, I photograph the scene and show you the picture, the information that the cup has coffee in it is conveyed in analog form. The picture tells you that there is some coffee in the cup by telling you, roughly, how much coffee is in the cup, the shape, the size, and the color of the cup, and so on.

On the other hand, their rich, multi-dimensional specificity also makes pictures computationally expensive: in order to represent anything at all, a picture must represent a lot, and in highly nuanced detail.

Sentential systems lie at the other extreme of reliance on direct isomorphism. First, their combinatorial and representational principles abandon any sort of direct physical isomorphism between vehicle and content. The semantic principles mapping the vehicle's constituents to represented contents are clearly highly arbitrary and conventional: neither the word 'tree' in English nor 'l'arbre' in French resembles a tree in any salient respect.[18] This arbitrariness frees sentential systems from any substantive constraints on the possible semantic values of their syntactic constituents. The syntactic principles combining those constituents are less arbitrary, but they too clearly abandon any appeal to physical isomorphism.[19] Instead, some sort of functional relation among syntactic constituents maps onto some sort of logical or metaphysical relation among the semantic values of those constituents; for instance, in the sentence

Socrates is wise

the syntactic relation of functional application mirrors a metaphysical relation of instantiation.[20] And in turn, both these syntactic relations and their logico-metaphysical counterparts can be embedded into indefinitely many further relations, to produce vehicles and correlative contents that are not merely indefinitely complex, but indefinitely hierarchically structured.

Note that it is only in this highly abstract sense that in sentential thought "the syntactic structure of mental states mirrors the semantic relations among their intentional objects," as Fodor and Pylyshyn (op. cit) claim that mental representations must do. The syntactic structure of pictures—and as we will see, of maps—mirrors semantic relations among the represented objects quite directly: a syntactic constituent's being next to or above another constituent in a picture or map mirrors the relation of proximity or aboveness among the represented objects or properties in the world. If we only attend to familiar natural languages, it can seem to be a deep requirement on thought per se that it must have subject/predicate structure, in order to mirror the deep metaphysical relation of objects possessing properties (e.g. Strawson 1963).[21] But for other cognitive purposes and given other representational formats, the distinction between individuals and properties may be comparatively marginal: perhaps *Socrates* is just a relatively stable property, or a comparatively homeostatic collection of properties, which can only be instantiated in one location at any given moment (cf. Quine 1960, Millikan 1998).

Because sentential systems represent by combining discrete, conventional symbols in an abstract structure, they are highly digital: they deliver chunks of information about discrete states of affairs. They also have a very minimal lower bound of informational content: a sentence can represent just *that* there is a cup, or *that* something is red, while remaining silent about every other aspect of the world. These features make sentential systems a computationally cheap means for tracking and categorizing information in small bits and at various levels of abstraction.

Cartographic systems are a little like pictures and a little like sentences. Like pictures, maps represent by exploiting isomorphisms between the physical

Kelly	Luke	Abraham
Dante	Lucy	Janelle

Figure 2. A map consisting entirely of words.

properties of vehicle and content. But maps abstract away from much of the detail that encumbers pictorial systems. Where pictures are isomorphic to their represented contents along multiple dimensions, maps only exploit an isomorphism of spatial structure: on most maps, distance in the vehicle corresponds, up to a scaling factor, to distance in the world.[22] Further, typically this spatial isomorphism itself only captures functionally salient features of the represented domain: for a road map, say, only streets and buildings and not trees and benches. Maps also depart from the direct replication of visual appearance by employing a disengaged, 'God's eye' perspective instead of an embedded point of view.

Where the syntactic principle that combines constituents in maps relies on a fairly direct, albeit selective, isomorphism, the semantic principle which maps those constituents to objects and properties in the world can be quite indirect and arbitrary. Road maps often represent churches with a cross, four-lane highways with a red line, state capitals with a star, and cities by their names. This further reduces maps' computational and informational load: rather than specifying the shape, color, relative size and orientation of a church, a map employs a minimal, easily replicable symbol to represent *that* there's a church at the relevant location. It also significantly increases maps' expressive range, by freeing them from the constraint of representing only visually perceptible features: for instance, an '×' on a pirate's map might represent buried treasure.

Indeed, some maps employ exclusively arbitrary, linguistic icons. For instance, the configuration in Figure 2 might constitute a map of students' assigned seats. Such a configuration is still a map, rather than a sentence, because it deploys the basic combinatorial principle of spatial isomorphism. Thus, when it occurs in a cartographic system, the icon 'Janelle' has the same function as every other constituent on the map: to indicate the relevant object/property's location in a spatial configuration alongside other represented objects/properties. By contrast, when 'Janelle' occurs in a sentential system, its syntactic function is different: it names an individual, and can only combine with expressions of appropriate functional types in a hierarchical structure. It is a notable feature of humans' representational abilities that they are sufficiently flexible to deploy the same expression in such different contexts.

Seating charts lie at the extreme of conventionalization; on most maps the constituent icons do share some salient resemblance to the objects and properties they represent. In particular, the physical features of the icons themselves often reflect salient physical features of the objects or properties being represented. Thus, a straight line represents a straight street and a crooked line a crooked street; a blue blob represents a pond of that very shape and exploits similarity of color to indicate that it's water; and a green blob represents a park of that very shape and exploits similarity in color to indicate that it's filled with vegetation. Thus, although maps employ discrete syntactic constituents with a significantly conventionalized semantics, there's still a significant interaction between their formal properties and mode of combination and what they represent. Nonetheless, the only strong constraint on the icons employed by cartographic systems, and on their potential semantic values, is that the icons' own physical features can't conflict with the principle of spatial isomorphism. Thus, one can't represent a street with a circle, not because it would be too arbitrary, but because this would make it impossible to place the icon in a spatial configuration that reflects the spatial structure of the represented content: for instance, one couldn't depict two streets as parallel, or as intersecting.

Other representational systems balance direct resemblance and abstract conventionality in different ways. On the one hand, pictographic languages combine an abstract, sentential syntax with a semantics that relies on visual similarity. On the other hand, diagrammatic systems, such as Venn diagrams, EKG charts, and bar graphs, don't necessarily exploit any physical resemblance in their semantics: the relation *being a child of* doesn't look like a line on a family tree. The principles by which their syntactic constituents are combined, though, fall interestingly between those of pictures or maps and those of sentences. Where pictorial and cartographic syntaxes use concrete spatial structure to represent concrete spatial structure, and where sentential syntax use abstract, functional structure to represent abstract, logico-metaphyical structure, diagrammatic systems often use concrete spatial structure to represent highly abstract structure. Thus, a Venn diagram might use intersections among circles to represent intersections among sets, while a bar graph might use height to represent annual expenditures. I diagram some of these interactions among representational systems in Figure 3.

There's obviously much more to be said about the syntactic and semantic principles that govern various representational systems, and about whether and how to draw boundaries between these systems. The crucial point for our purposes is just that many maps employ discrete, recurring constituents with a highly arbitrary semantics, and combine them according to systematic rules to produce systematically related whole representations. But at the same time, the principle according to which those constituents are combined relies on a spatial rather than purely logical isomorphism between the structure of those constituents and the structure of the corresponding elements in the content. This demonstrates in concrete terms that there is more than one way in which "the syntactic structure of mental states [can] mirror the semantic relations among their intentional objects," as Fodor et al. take the argument from systematicity to require. As a result, Premise 4 in the Argument for the Language of Thought is simply false: there are non-linguistic combinatorial representational systems.

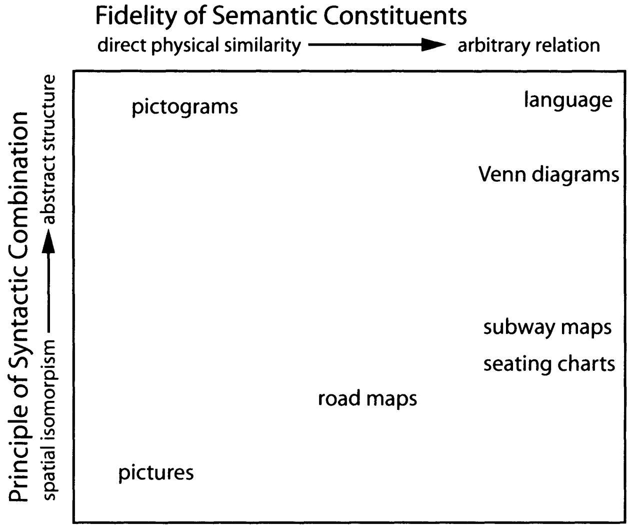

Figure 3.

§3: The Representational Advantages and Disadvantages of Cartographic and Sentential Systems

In §2, I established that both sentential and cartographic systems employ recurrent constituents combined according to systematic rules, but that their compositional principles differ significantly. This demonstrates that maps aren't just languages written in a funny notation, and hence that they constitute a potential counterexample to Strong-LOT. However, to demonstrate the falsity of Strong-LOT, we also need to show that a non-sentential system could fulfill the basic cognitive functions of representing and reasoning. In this section, I argue that so long as a thinker's representational needs are sufficiently simple, it could think largely or entirely in maps; indeed, in important respects a cartographic system would be easier for such a thinker to use. However, as the range and complexity of contents a thinker needs to represent and reason about increases, maps become increasingly cumbersome. This gives us good reason to think that much of our own thinking does occur in sentences.

One reason it seems implausible that a thinker could do all or even most of its thinking in pictures, besides heavy computational demand, is that pictures' high semantic density and syntactic complexity makes it hard to see how one could use them to reason: many of the changes one can make to a picture will destroy its structural coherence. By contrast, because maps employ discrete icons with a potentially conventionalized semantics, and abstract away from so much detail, they have a significantly wider expressive range and permit considerably more flexible manipulation. A thinker can easily place, remove, and relocate a wide range of symbols on a map without destroying the rest of the map's structural coherence. At the same time, though, maps still share with pictures the ability to present lots of information simultaneously in a compact way. This combination of features makes maps especially efficient vehicles for certain kinds of reasoning.

In particular, maps automatically conjoin information about the spatial locations of all the objects and properties they represent. Thus, suppose I have the following sentences specifying the locations of Bob, Ted, and Alice:

Bob is at the grocery store at 10th and South.

Alice is at the café at 11th and Pine.

Ted is at the park at 9th and Spruce.

9th Street is east of 10th Street.

10th Street is east of 11th Street.

Lombard Street is north of South Street.

Pine Street is south of Spruce Street.

Lombard Street is between South and Pine Streets.

Faced with these sentences, I still have to do considerable cognitive work to figure out how Bob, Ted, and Alice are located in relation to one another. By contrast, if Bob, Ted and Alice are each represented on a map, as in Figure 4, then not just their respective locations but also the relations among them are explicitly represented and cognitively transparent.

This point is familiar to anyone who has taken the SAT, GRE, or LSAT, which often include word problems requiring one to deduce the relative locations of various objects. By far the most efficient way to solve such problems is by constructing a map, and all of the cognitive work comes in that construction: the solution is automatically available once the information has been encoded. As Shimojima (1996) puts it, the 'inference' from premises to conclusion comes along as a "free ride".[23]

An important corollary of this is that maps are holistic representational systems, while sentential systems are atomistic (cf. Braddon-Mitchell and Jackson 1996, 171). In Figure 4, no single, syntactically isolated portion of the map represents just where Bob is, without also representing Bob's location relative to Ted and Alice and everything else that is represented on the map. Each icon contributes to the overall spatial configuration, and the location of each object and property is given in terms of this overall configuration. As a result, any alteration in the location of the 'Bob' icon automatically alters the represented relations between Bob and everyone else.[24]

This difference in how sentential and cartographic systems conjoin or accumulate information means that they are likely to distribute the task of

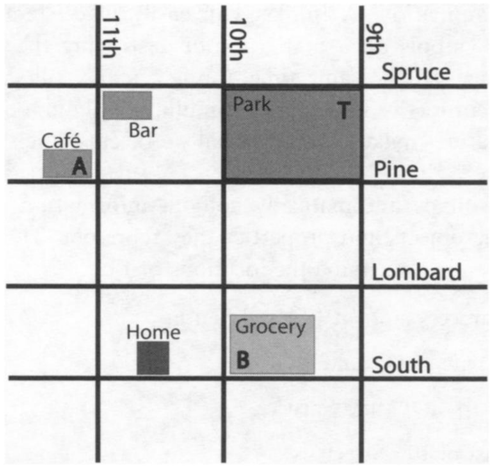

Figure 4. A map representing multiple locations in relation.

representing the same overall information in very different ways. A map can easily represent the locations of and relations among many objects and properties in an explicit yet cognitively transparent way, thereby minimizing the need for processing to recover those locations and relations. By contrast, it would be massively cumbersome to spell out this same information in sentences: a practically feasible sentential representation will only specify some of that information explicitly, and will rely on processing to make latent information explicit. However, because the number of further sentences one can derive from any substantive set of initial premises is so large, it's not feasible to just crank out that information by brute force. For practical purposes, a thinker needs a content- and context-sensitive way to extract relevant information. Thus, when dealing with relative spatial locations, sentential systems face a processing challenge, and a risk of processing error, that cartographic systems don't.

So long as a thinker works with a single map, she has neither need nor room for an explicit representation or process of conjunction: the map itself has already taken care of it. A thinker might also operate with multiple maps, though, which won't automatically accumulate their respective information into an integrated whole. Such a thinker would thus need some way to collate their information. If the maps represent sufficiently continuous regions of space, then conjunction can proceed by concatenation and superimposition, controlling for scale and orientation. However, maps representing spatially discontinuous regions cannot be syntactically conjoined. The contents of discontinuous maps may still be related, though, in ways a thinker needs to be sensitive to: for instance, two maps of distinct spatial regions might be inconsistent if they both represent Bob, but not if they both represent a Dunkin' Donuts. These higher-order relations between maps can be captured in implicit rules for using the maps; but they can't themselves be represented explicitly on any map. At most, the system can employ a symbol like '&,' which itself lacks any spatial significance, to connect distinct maps. This differs markedly from sentential systems, where conjunction can be explicitly represented in a fully general way.

Negation, Disjunction, and If-Then

While normal maps and their cognitive analogues are significantly more efficient than sentences at conjoining information about related spatial locations, such maps lack any means to explicitly represent the other truth-functional relations. This is a significant limitation in expressive power, to say the least. First, consider negation. On the familiar maps we ordinarily use to navigate—say, a Rand-McNally map of Philadelphia—the absence of an icon from a point on the map represents the absence of the correlative object/property from the correlative location in the world (cf. Rescorla 2005). However, this is an artifact of our treating Rand-McNally maps as omniscient with respect to the total presence and absence of any type of property or object they represent. By contrast, an ordinary thinker constructing her own cognitive map obviously won't be omniscient, and so can't employ our ordinary interpretive rule. Such a thinker will still likely need to keep track of negative information, though: say, that Bob isn't home, or that Alice isn't at the store.

In principle, it's not hard to extend maps to represent negative information. Most crudely, we could introduce a higher-order icon with the force of a 'contrary operator': say, putting a slashed circle over the 'Bob' icon to indicate that Bob is not at the represented location. Because we are already employing symbolic icons as constituents, this doesn't itself fundamentally change the sort of representational system we're employing. However, this technique would quickly lead to massive clutter. A more elegant solution would color icons and background regions to reflect positive and negative information.[25] For instance, the default state could be a grey background, expressing neutrality about the presence and absence of every potentially representable object and property. A black (or other fully-saturated) icon would represent certainty that the relevant object/property is at that location, while a white (or anti-colored) icon would represent certainty of its absence; a white background could then represent certainty that there were no other, unrepresented objects or properties in that region besides those explicitly represented on the map. Intermediate values for coloring icons and backgrounds could track finer variations in positive and negative credence.

Maps can also be extended, in principle, to deal with disjunction and if-then. Because maps work by placing discrete icons in determinate configurations, standard maps lack any way to represent partial information, such as that *either* Bob is at the store *or* he is at the bar; or that *if* Bob has gone to the store, *then* he's walking this way. It is possible, but inefficient, to represent disjunctive and conditional contents with non-spatial symbols relating distinct maps: when the maps significantly overlap, the system needs a way to isolate just their salient representational differences; and even when maps don't largely overlap, the relevant information being disjoined or conditionalized often won't encompass the entire content of either map, but just a selected element, such as Bob's location. In either case, in order to act on any disjunctive or conditional information it represents, the system needs some way to isolate specific elements within maps. The case of negation offers a better model: thus, one might color (sets of) icons with alternately flashing yellow lights to indicate that one or the other state obtains. Likewise, one might use solid blue lights to indicate the antecedent of a conditional, with flashing blue lights to indicate its consequent. This method could also be used represent other features that cannot be expressed in standard maps: for instance, one might represent past or future tenses by writing the icons in italics or cursive. Given these suggested extensions, Figure 5 might be used to represent, among other things, that Ted is not at the park, that Alice is either at the café or the bar (with staid dashed lines replacing flashing yellow lights), that Bob was at the grocery store and that no one else is, that no one is at home, and that the thinker herself is on Lombard west of 11th Street.[26]

Intensionality and Quantification

Two other important sorts of information are trickier to represent explicitly in maps, but can arguably still be managed, at least in principle. First, consider intentional attitudes. We can use the same basic method to represent some of

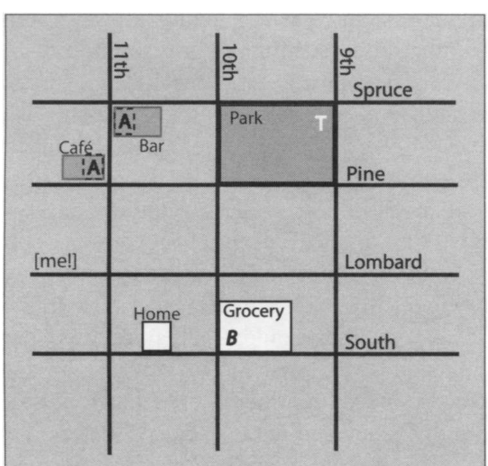

Figure 5. An extended map, representing negative and disjunctive information, and past tense.

the thinker's own mental states other than belief: the desire to be at the café, for instance, might be represented by placing the 'me!' icon over the 'café' icon with a flashing green light, and the fear that Bob is at the bar might be represented by placing the 'Bob' icon over the 'bar' icon with a flashing red light. Other agents' attitudes are more challenging. If we represent that John believes that Bob is at the bar simply by superimposing an icon for 'John believes' onto the 'Bob' icon at the 'bar' location, then we risk attributing to John the belief that Bob is also related to all of the other features represented by the map: after all, the map individuates the location that John is represented as believing Bob to be at in terms of that location's relation to the entire configuration. But John might not believe all or any of the rest of this configuration. Thus, we either need to keep the thinker's map of John's beliefs entirely separate from her own, or else to isolate specific regions within her own map as reflecting John's beliefs. The challenge for the first model is that the thinker may have little further information about how else John believes the world to be, and so will end up with mere map-snippets, which still need to be integrated with the thinker's own belief-map by means of merely implicit rules for use. The challenge for the second model is to respect the intensional quality of John's mental states without obliterating co-located information on the thinker's own map.[27]

Quantificational information poses the final and most serious challenge for cartographic systems.[28] Maps easily represent some sorts of existential information, such as that *there is a café here*. However, because maps work by placing determinate features at definite locations, they can't represent information that's not spatially located, such as that *somebody, somewhere is wearing a red shirt and carrying a gun*. As we might put it, the bare existential information that *something or other, somewhere or other, is F* falls below the minimum bound of cartographically representable information.[29] Specifically, because so many desires are for things the thinker can't locate—if she could, she'd go get them—it is especially hard for maps to capture the full range of desires that an agent is likely to have.

At the other end, universal information, such as that *all the Fs are G*, can be too 'big' to fit on a map. A map may have a *G* icon everywhere there's an *F* icon. But this leaves out precisely the fact we're interested in: that those are all of the *F*s. If a thinker treats the map as authoritative, then it does implicitly contain the information that all of the *F*s are *G*s, because the thinker could extract this information by checking every *F*, noting that it is always accompanied by a *G*, and noting that there are no other *F*s on the map. And perhaps one could introduce a (non-spatial) symbol to mark maps as authoritative. But as I said above, individual thinkers' maps are unlikely to be authoritative across the board. A thinker thus needs some selective means to represent *that* all of the *Fs* are represented, without also representing that all of the *J*s are. We could use the model of negation to do this, say by writing all of the icons of a given type in bold when the thinker believes that all instances of the relevant type are represented on the map. But this would still only permit the implicit representation of the information that *all the Fs are G*: the thinker would still need to extract the universal quantification by checking all of its instances individually. More importantly, a thinker might well believe that *all the Fs are G* without having any beliefs about the specific locations of the *F*s, and so without being able to place them at definite locations on the map.

Expressive Limitations and Usability

Because one can beef up ordinary maps in all of these ways, there are fewer absolute expressive limitations to cartographic systems than we might naïvely suppose. In particular, because maps exploit discrete, recurrent syntactic constituents with stable, at least partlyconventionalized semantic properties, one can achieve something close to the effect of sentential structure within a cartographic system by manipulating the basic icons in ways that don't affect their spatial structure. In effect, we've introduced rules for generating syntactically complex icons which represent semantically complex objects and properties: not-Bob, past-Bob, etc.[30] So long as these icons still function as labels placing objects and properties at locations, one might argue, and so long as their mode of combination sets up an isomorphism between their spatial structures and those of the analogous features in the world, we're still operating within a fundamentally cartographic system. What would fundamentally alter the nature of the representational system would be to assign some *other* representational significance to the spatial relations among icons—say, so that placing two icons next to each other sometimes meant that the correlative objects were near one another, but other times meant that the leftward one loved the rightward one. Likewise, one might insist that it would fundamentally alter the representational system if one employed a fully sentential syntax to combine the icons into a complete representational unit in their own right, so that placing the icons for *dog*, *bit*, and *boy* together on the map represented that at that location the boy bit the dog.[31] If we decided to take this limitation seriously, then although one could legitimately introduce the sorts of icons we've discussed for higher-order relations like *if-then*, and although one could introduce icons for properties like *food,* and icons for types of objects like *happy guy,* one couldn't legitimately introduce icons with predicative force, to represent properties like *being happy*, *being bald*, or *loving*.

I'm less interested in drawing sharp boundaries between types of representational systems than in getting clear about how typical instances of each system work, and on the implications of this for what they, as well as hybrid systems, can represent and how one reasons with them. If we're really interested in boundaries, perhaps we should rule out the use of fully conventional symbols like 'Janelle' or 'Philadelphia', and consider only topographical maps without words to be 'real' maps. As I emphasized in §2, though, maps are themselves an interesting hybrid between the direct replication of visual appearance employed by pictures and the fully abstract, conventional representation employed by sentences. More generally, hybrid systems are often so useful and elegant precisely because they synthesize the expressive advantages of distinct representational systems (cf. Tufte 1990).

The more important point concerns how the extended systems can be used to represent and reason. So long as a thinker is merely representing objects' and properties' relative spatial locations, maps' holistic, accumulative quality makes them efficient representational vehicles compared to the cumbersome atomistic representation of sentences. But as we extend them to accommodate the representation of more complex contents, maps become much more unwieldy.

First, at a practical level, extended maps are harder to use. No representational system can make its vehicles fully cognitively transparent: even the most 'obvious' representational system still requires some background knowledge in order to use its vehicles appropriately. With pictures, a user must know to treat realistic pictures as replicating visual appearances, and to treat impressionistic and cubist pictures as distorting or filtering visual appearance in certain ways. For standard maps, the user must know to treat the map as spatially isomorphic to a specific region in the world, subject to an orientation and scaling factor; she must know the semantic significance of the constituent icons; and she must know whether to treat the map as authoritative. Even so, standard maps are *comparatively* cognitively transparent. They always and only employ icons to represent objects and properties as arranged in a spatial configuration, and they represent this configuration by replicating that very same configuration among the icons themselves. As a result, if a thinker can locate herself on the map and orient it to reflect her current orientation, she can navigate in the world by moving in the very same direction as, and by a distance that is directly proportional to, the direction and distance in the map. By contrast, the extended system exemplified in Figure 5 requires the user to employ a variety of interpretive principles, many of them quite abstract. Further, even if we exploit dynamic features like flashing lights rather than clunky lines, such a system will inevitably have considerably more cluttered vehicles, and be more prone to encoding and processing errors. By contrast, although sentential systems have high 'entry costs', once the basic syntactic principles have been mastered it is quite easy to construct and understand sentences of indefinite complexity about a wide variety of contents.

Second, although we haven't identified any absolute in principle barriers on kinds of information that maps can be extended to represent, there are quite serious limitations on the full generality of their expressive range. Where the syntactic operations by which a sentential system represents conjunction, negation, disjunction, conditionalization are fully general and easily executed, even an extended map only permits the explicit conjunction, disjunction, and conditionalization of bits of information that are spatially related. Likewise, in sentential systems it is easy to selectively represent one abstract state of affairs while remaining neutral about the particular concrete facts in virtue of which it obtains. By contrast, a map can only represent an abstract state of affairs by specifying the locations of the underlying objects and properties which make it true. Specifically, a map can only represent predicative, intentional, and quantificational information by placing icons for particular objects and properties at particular locations. But a thinker may not be able to locate all of the relevant objects and properties, and may have no immediate cognitive need to do so.

Third, where the recursive structure of sentential systems makes it easy to represent hierarchically-structured contents of indefinite complexity, the extensions I've suggested to cartographic systems all operate at the same level, applying directly to the basic icons, or at most to collections of such icons, which themselves all serve to place objects/properties at locations. In principle, the system might be extended even further to permit those higher-order relations to apply to one another with different scopes. But in effect, this will require importing the full hierarchical recursivity of language into the cartographic system—all without interfering with the basic principle of spatial isomorphism. I've talked about fonts, background colors, and flashing and solid colored lights in order to provide some concrete sense for how an extended cartographic system might represent higher-order relations by non-spatial means. But there are only so many non-spatial but still physical ways to manipulate icons. To represent multiply embedded higher-order relations, and to represent multiple higher-order relations of the same kind on a single map, we will eventually need something like sentential notation.

Together, these points about the generality, selectivity, and indefinite hierarchical structure of sentential systems make sentential systems much more efficient vehicles for the representation of abstract, complex, hierarchically-structured information. By contrast, even if a cartographic system can be extended in principle to express such information, that representation will be massively cumbersome. Thus, suppose that a map is capable of representing the contents that *some of the ballerinas are at the bar,* that *some of the ballerinas are at home*, and that *all of the officers are at the bar*; suppose also that we have a way to represent past tense, and that we permit the introduction of an icon for the relation of dancing. In principle, this should enable the map to represent the information that *some but not all of the ballerinas danced with all of the officers.* However, it will be vastly simpler to express this content in sentential form—let alone to represent the content that *if some but not all of the ballerinas danced with all of the officers, then no ballerina is both tired and jealous*, or even more complex contents.

Thus, the original source of maps' representational strength—their use of direct spatial isomorphism—is ultimately also the source of their representational weakness. Cartographic systems are sufficiently systematic to satisfy the basic requirements of representation and reasoning that motivate the arguments for Weak-LOT. And because they employ a different combinatorial principle than sentences, they demonstrate the falsity of Strong-LOT. But because the basic combinatorial principle of maps, as of pictures, relies on a direct isomorphism between physical properties of the vehicle and those of the represented content, maps are significantly less flexible than sentential systems. To achieve a really robust expressive richness, which is capable of selectively representing and fluidly manipulating the sorts of abstract, hierarchically-structured contents that make human cognition so distinctively powerful, a representational system needs to employ syntactic and semantic principles that are sufficiently abstract that they don't themselves impose any substantial limitations, either on what semantic values can be assigned to the syntactic constituents or on the complexity with which those constituents can be combined.

Diagrammatic systems come considerably closer to this ideal, because they are free to employ an isomorphism between the vehicle's physical structure and *any* sort of structure, including logical and metaphysical structure, in the represented content. This makes them very useful for representing and reasoning about abstract and hierarchically-structured information in a way that is still comparatively intuitive because it still exploits basic geometry. In particular, Venn diagrams are useful for reasoning about quantificational information; and family trees and flow charts can represent social and causal relations in a compact, obvious way. Indeed, given that we often employ diagrams to illustrate the logical structures of sentences, there may ultimately be no principled boundary between diagrammatic and sentential systems. However, most diagrammatic systems are considerably more restricted than full-blown languages, because they assign a dedicated interpretation to the vehicle's topological structure: in the case of family trees, say, the ancestor-descendant relation. More importantly, precisely because their syntax still exploits physical properties of the vehicle, many diagrammatic systems face significant expressive limitations of their own: for instance, some Venn diagrams involving four or more circles cannot be drawn in a single plane (Lemon and Pratt 1998).

We thus arrive at a position we might call *Sophisticated-LOT*: the representational vehicle which underwrites highly flexible thought about abstract, hierarchically-structured states of affairs is likely to be sentential in form. Because the distinctive power of human cognition seems to depend on our agility at representing and manipulating such contents, this gives us good reason to think that much of our own cognition, in contrast to that of other animals, takes place in language. However, this conclusion depends crucially upon the specific contents that humans think about and what they do with those contents, and not on general features of thought *per se*.

§4: Does It Matter?

At this point, we've seen that both sentential and cartographic systems employ discrete, recurrent parts and systematic combinatorial rules to represent systematically related contents; but also that they employ importantly different combinatorial principles, and hence that the two systems differ in how they can be used to represent and reason about states of affairs in the world. We are thus now in a position to respond more fully to the objection that the distinction between Weak- and Strong-LOT is not theoretically significant.

The Objection from Informational Equivalence

A dismissive attitude toward the difference between Weak- and Strong-LOT might seem especially warranted if we take seriously Pylyshyn's point, cited at the end of §1, that the fundamental unit of analysis must be the entire package of representational vehicle and rules for use. Different syntactic structures clearly require different formal transformations; but perhaps all that matters is that there be *some* rule-governed and reliably truth-preserving way to go from representational inputs into outputs. John Anderson (1978, 262-3) supports this attitude in the context of the mental imagery debate by arguing that any behavioral data can always be accounted for by either a sentential or an imagistic representational system, because the two systems will distribute the representational labor differently between vehicle and process:

> [I]t is not possible for behavioral data to uniquely decide issues of internal representation.... One can show that given a set of assumptions about an image representation and a set of processes that operate on it, one can construct an equivalent set of assumptions about a propositional representation and its processes. Or one can be given a propositional theory and construct an equivalent imagery theory. In fact,... given any representation-process pair, it is possible to construct other pairs with different representations whose behavior is equivalent to it. These pairs make up for differences in representation by assuming compensating differences in the process.

Likewise, Randy Gallistel (1989, 172) concludes that behavioral data can't distinguish between the hypotheses that bees represent the world by maps or by "the equivalent of a surveyor's field notes":

> Since the information content of the surveyor's notes and a cartographic product based on those notes are the same, it is going to be difficult to decide unequivocally from behavioral work alone what actually occurs inside the bee's nervous system.

First, it's obviously true that two systems may be "informationally equivalent" in the sense that in principle one can extract the same information from each system, or that they make the same 'cut' in the space of possible worlds. However, this notion of "information equivalence" is highly rarified. As we saw in §3, in practice plausible cartographic and sentential systems will distribute the representational burden between vehicle and process very differently for different contents, with sentential systems relying much more heavily on processing to recover implicit information about spatial relations, and cartographic systems relying on processing to recover implicit information about universal quantification. So far, this just illustrates Anderson's claim. But this difference in representational burden in turn means that in practice, the same bit of information may be easily accessible in one system and recoverable only with much effort in the other. This point extends even to pairs of sentential systems that differ in which premises they store explicitly and which algorithms they employ to recover latent information. Thus, given the actual temporal and processing constraints on practical decision-making, we should expect that one representational system will fail to recover in-principle-represented-but-latent information where the other succeeds, despite their overall in principle informational equivalence.

Second, in practice we are also likely to be able to distinguish informationally equivalent vehicle/process pairs behaviorally by observing differences in the two systems' 'failure modes' (cf. Marr 1982). In particular, the fact that maps are holistic forms of representation while sentences are atomistic means that each system is likely to break down in a distinctive way. Thus, to the extent that a thinker fails to exploit the full consequences of information acquired on distinct occasions to achieve her goals—for instance, if a rat undertakes separate trips to get water and food, returning to its nest in between, although it would be shorter to go directly from the water to the food—we have some evidence that it stores information atomistically. Conversely, to the extent that a thinker automatically integrates information from separate experiences, this supports the hypothesis that it employs a more holistic system. For instance, if a bee regularly sets out on the most efficient route home when released in a new spot, or if one illusory experience ramifies error throughout the thinker's behavior, or if disorientation prevents a rat from taking any sort of action, then this gives us some reason to believe that the thinker is employing something like a cognitive map. Any by and large, empirical evidence about the navigational skills of rats, bees, and other animals does support the claim that they often employ some sort of map-like system (cf. e.g. Boesch and Boesch 1984, Gould 1986, Gallistel 1998).

Of course, no single piece of behavioral evidence, or even any particular collection of evidence, can be absolutely dispositive here. A thinker *might* reliably display behavior suggestive of a holistic representational system because she's extremely good at deducing consequences from sentential premises. Alternatively, a thinker *might* fail to exhibit cognitive closure because she represents the world with multiple distinct maps and has failed to compile the information on them. Still, in a practical context we should expect the different ways that sentential and cartographic representational systems are likely to distribute information between vehicle and process to manifest themselves behaviorally. Thus, the fact that two representational systems are informationally equivalent in principle doesn't show that there can't be any significant empirical justification for claiming that a thinker is employing one rather than the other system.

The Objection from Expressive Equivalence

A more hard-line version of the objection from informational equivalence insists that the only difference between representational systems that really matters is expressive power. So long as a theorist's preferred representational format has the power to represent the contents he attributes to a thinker, no behavioral evidence can force him to abandon the hypothesis that the thinker is employing that format. This objection is typically advanced by 'propositionalists' like Pylyshyn (e.g. 1973, 2002, 2003) against the claim that a thinker is employing a pictorial or cartographic system: if everything that can be expressed in maps or pictures can also be expressed in sentences, what could ever demonstrate that a thinker is using pictures or maps instead? The objection seems especially forceful given that many philosophers endorse the idea that language is expressively complete (cf. e.g. Searle 1969).

In §3, we saw that in principle, cartographic systems are less expressively limited than one might naïvely suppose. Thus, so long as there's no direct evidence that a thinker is representing, say, non-localized quantificational information, a pro-cartographic theorist might doggedly insist that all the evidence about what a thinker represents is compatible with the hypothesis that she is thinking with maps.[32] Further, in principle even sentential systems have expressive limitations. As we saw in §2, sentential systems are highly digital: they combine discrete, arbitrary symbols in an abstract hierarchical structure. But this in turn means that at any given moment, a given sentential system only has the expressive resources to represent countably many contents: those formed by all the combinations of its syntactic constituents. By contrast, because pictures and maps are analog modes of representation, they are potentially continuous; and as such, they can represent continuously many contents. In particular, a cartographic system with the expressive resources to draw continuously differentiated blobs of the sort in Figure 1 already contains within itself the expressive resources to represent ponds of uncountably many shapes, as well as to configure the various types of objects and properties it can represent in uncountably many ways. It is true that a sentential system can typically expand its vocabulary, either by directly ostending new features in the world, or else by exploiting a systematic isomorphism between new expressions and features in the world: for instance, by naming shades of blue 'Blue 100', 'Blue 101', etc., where each consecutive shade is just discernably more saturated than the previous one. However, both of these methods are themselves dependent on conditions that may not always be met. A thinker can only introduce expressions by ostension for those features that she can actually ostend; but sometimes she may lack the appropriate cognitive or causal access to those features (Camp 2006). Likewise, a thinker may lack any systematic way to match new expressions to properties by means of systematic isomorphism of the sort envisioned for color. More importantly, these methods still only expand the vocabulary in a countable way, and so don't enable the language to represent continually many states of affairs.

These considerations about expressive power and limitation are highly theoretical, of course. In practice, I do think considerations of expressive power strongly favor sentential systems: they do a much better job at expressing a much wider range of contents, especially highly complex, abstract, hierarchically-structured contents. Further, in practice it's unlikely that any representational vehicle for thought will be truly continuous, or that a creature will need to represent uncountably many contents. At the same time, though, a creature who lives in a messy world with features that differ saliently in a fairly continuous way will need a fairly continuous means to represent those features. Thus, even if a sentential system is capable of encoding the relevant information, it will be much more useful to employ a format, such as a diagram or map, which directly represents fine-grained differences along one dimension while abstracting from detail along other dimensions. Thus, the question of which representational format best reflects a thinker's representational needs, abilities, and limitations—and with it, the question of which format it is most plausible to assume a thinker is employing—depends largely on what sorts of contents the thinker most often represents, and how she needs to manipulate them.

But What About the Brain?

The final, and most pressing, objection I want to consider attacks the possibility of non-sentential thought from a more empirical angle. The central point of my discussion of cartographic and pictorial systems has been that they employ a principle of spatial isomorphism between vehicle and content. And this obviously entails that pictures and maps themselves *have* a spatial structure. We know what this means for a normal physical map—the kind that's written on paper or built with twigs and twine, and those are the terms in which I've been discussing the relative expressive powers and limitations of maps and sentences. But how are we to interpret this claim in the context of thought?[33] If the claim that thinkers employ cognitive maps is read as the claim that there are spatial structures *in the brain* isomorphic to spatial structures in the world, the objection goes, then this is radically implausible. By contrast, precisely because sentential systems employ such abstract semantic and combinatorial principles, the claim that a thinker employs a language of thought is compatible with an extremely wide range of plausible neurological implementations. Thus, although by itself Weak-LOT leaves open the possibility of thought with a non-sentential form, one might think, only Strong-LOT offers an empirically plausible implementation of Weak-LOT.[34]

This objection raises issues about neural processing that are beyond the scope of this paper. In response, however, note first that physically instantated maps are in fact ubiquitous in the brain. Scientists have known since the 1940's that the mammalian cortex represents many aspects of the world, especially the layout of one's own body and sensory stimuli, in such a way that the spatial structure of neural firing reflects the physical or psychological structure of the represented content.[35] There is also evidence that more abstract information, such as the relations among keys in Western tonal music, can be represented topologically.[36] Thus, the objection can't depend on a wholesale rejection of spatial isomorphism in cognitive representations. Rather, the objection must be something more like the worry that it's implausible that the brain contains enough room for the entire world to be represented cartographically. This objection does have bite against pictures, because they are so computationally expensive. But maps abstract away from much of pictures' detail, and are free to employ highly abstract icons. Thus, the computational demands imposed by maps are potentially much closer to those for sentential systems than to those for pictures.

More importantly, the claim that thought might be map-like rather than sentence-like is best interpreted functionally. As Fodor and his co-authors themselves emphasize, the Language of Thought hypothesis is not committed to any particular claim about the particular neural instantiation of cognition; indeed, computationalism is compatible with connectionism at the level of actual neural firing.[37] Rather, the Strong Language of Thought hypothesis is the claim that at "the cognitive level of description," an adequate causal account of a system of mental representations and reasoning must type neural processes in terms of word-like constituents and language-like rules for combining them. But this same interpretation is available to someone who claims that at least some thinkers, such as bees and rats, employ a cartographic system for thought. Fodor and Pylyshyn (1988, 13) claim that they take claims about combinatorial structure "quite literally" insofar as they assume that

> the combinatorial structure of a representation is supposed to have a counterpart in structural relations among physical properties of the brain. For example, the relation 'part of', which holds between a relatively simple symbol and a more complex one, is assumed to correspond to some physical relation among brain states.

They do not take LOT to require that the first physical state actually be a part of the second one. Likewise, the cartographic theorist can take claims about cartographic structure quite seriously, if not fully literally, by maintaining that relations like 'next to', 'above', and 'intersecting', which hold between symbols in a map, correspond to *some* physical relations—not necessarily spatial—among brain states.[38] Given that we have identified substantive differences in how maps and sentences represent their contents, and in the patterns of reasoning that thinkers using them will employ, we can get some significant grip at the purely functional level on which format a thinker is employing. There's no reason to think that this cognitive level must itself be underwritten by another functional level at which the syntax and semantics are specified sententially.

§6: Conclusion

Throughout this paper, I have been operating with the fiction that a thinker only employs a single representational format for thought. I've done this in order to demonstrate in the starkest possible terms the falsity of Premise 4 in the argument for a Language of Thought, and so the falsity of the claim that thought *per se* must be language-like. If a thinker's representational needs are sufficiently restricted, then a wholly cartographic system could serve as a feasible vehicle for its thought. Indeed, if you're designing a cognitive system whose primary challenge is to navigate a fairly stable terrain in search of only moderately mobile features—food, water, shelter, and a mate—then maps provide an exceptionally efficient and computationally tractable system for representing and reasoning about the world. The limitations for maps lie in their inability to represent highly complex, hierarchically structured and abstract information in a fully general, selective, and flexibly manipulable way. Sentential—and diagrammatic—systems are considerably better equipped to handle these sorts of contents.

Ultimately, any plausible cognitive system, including especially our own, is likely to be highly multi-modal: storing and manipulating information in the formats of multiple sensory modalities, and centralizing information in cartographic, diagrammatic, and sentential formats. If this is right, of course, then it becomes commensurately more difficult to make principled predictions about what a thinker should or shouldn't be able to do. In one way, this simply illustrates the dangers of attempting to decide on entirely theoretical grounds what the form of thought must be. However, I also think that the differences we've identified in *how* pictures, maps and sentences work, and in turn in *what* they are good at representing, are substantial enough to suggest experimental means for testing which representational format a thinker actually employs when faced with specific cognitive challenges.

Notes

*Thanks to members of the Philosophy Department at the University of California, Davis, especially Adam Sennet and Paul Teller; to the Corridor Reading Group, especially John Collins and Ted Sider; and to Dmitri Tymoczko for extensive discussion.

1. That is, you can think other thoughts about the same objects, properties and relations under the same mode of presentation. Fodor, Evans, and other advocates of this argument generally individuate contents in Fregean terms, so that the content or object of thought is an abstract proposition, which is about a state of affairs in the world by way of a mediating sense. I prefer to individuate contents in Russellian terms, and to individuate thoughts and representational abilities in terms of contents plus modes of presentation; but nothing turns on this in what follows.
2. By contrast, one might claim there aren't any psychological laws about connections between thoughts that don't reflect some relation in their contents.
3. Crane's allusion is to the following passage from Frege (in a letter to Peano): "Where inferences are to be drawn . . . it is essential that the same expression should occur in two propositions and should have exactly the same meaning in both cases. It must therefore have a meaning of its own, independent of the other parts of the proposition."

4. Indeed, it seems likely that their representations are further structured, into (at least) a rank ordering of dominance relations.
5. Some proponents of a Language of Thought think of it as a largely innate Mentalese (Fodor 1975), while other think of it as an internalized natural language (Carruthers 2002). For the purposes of this paper, the difference between these is not significant.
6. Although it appears that Crane and Bermudez are disagreeing in the passages I've cited, both are spelling out the consequences of Frege's point, cited in fn. 3, that inference requires common constituents with stable meanings. The appearance of disagreement arises because Bermudez means something closer to Russellian propositions by 'content'.
7. Although Bermudez's talk of formal rules might appear to suggest that only logically valid inferences are ultimately justified, by itself the formality condition doesn't specify which formal rules a thinker should employ. One could have a rule of 'inductive inference' which specified that sixteen instances of *An F is G* and no instances of *An F is not G* entitled one to conclude *All Fs are Gs*. This is undoubtably a silly rule, but relative to an inferential system which licensed such a rule, the relevant inference would count as justified.
8. Bermudez and other rationalist philosophers, such as Sellars (1956/1997), Davidson (1975, 1978/80), and McDowell (1994), also hold that a transition between thoughts only counts as reasoning if the thinker is capable of higher-order reflection upon the transition's epistemic credentials. But such meta-representational reflection, they argue, is only possible in language. I think we should reject a requirement of higher-order reflection (Camp forthcoming). Humans have been solving complex problems for much longer than they have been providing formal justifications for their solutions. As Huw Price (1990, 231) observes: "It is tempting to think that agronomy is an older profession than epistemology—that we could think talk and argue about grass (and many other important things) at least an evolutionary step or two before we hit on the concepts of warrant, reason, justification, and the like." Further, not everyone agrees that higher-order thought requires language; see Origgi and Sperber (2000).
9. Analogously, one might deny Evans' claim that our ordinary understanding of thought commits us to thinkers' having fully general representational abilities. Indeed, Evans himself admits that full generality is an "ideal, to which our actual system of thoughts only approximately conforms" (1982, 105); further, he imposes categorial restrictions on the Generality Constraint to avoid the conclusion that failure to grasp the thought that *Caesar is a prime number* impugns one's competence with the concept *being a prime number*. I reject such restrictions on the range of concepts' application (Camp 2004), but I agree that full systematicity is an ideal to which actual thought only partially approximates (Camp forthcoming). Finally, one might challenge the claim that *language* is so systematic, either syntactically or semantically; see e.g. Johnson (2004) and Travis (1994).
10. Faced with such apparent failures of systematicity, theorists often postulate distinct, encapsulated modules within which systematicity is preserved (e.g. Cosmides and Tooby 1992, Carruthers 2002). Whether or not this is the appropriate theoretical response, it represents an important departure from the model of full systematicity.

11. Indeed, sometimes the only available explanations for a transition's validity will appeal to content rather than form, as in the case of inferences that combine sentences with other representations, such as maps or diagrams (cf. Barwise and Etchemendy 1990).
12. As Davies (somewhat disingenuously) comments on the passage from Evans above: "we can be in total agreement with Evans.... It is certainly no part of the LOT hypothesis, as it is argued for here, that the conscious, thinking subject is presented with thoughts as entities with non-semantic properties" (1991, 245).
13. As Fodor and Pylyshyn (1988, 32) put it, that "mental representations characteristically exhibit a combinatorial constituent structure and a combinatorial semantics."
14. In ch. 8 of his (2003), Bermudez does argue that only language enables metarepresentation, and with it "the inferential evaluation characteristic of second-order cognitive dynamics" (2003, 162). However, he assumes without further argument that only a linguistic vehicle can formally reflect inferential relations among contents.
15. Given that Fodor (1975) coined the phrase "Language of Thought," and that he and his co-authors so frequently move directly from talk about structured vehicles to a language of thought, it's natural to assume that they also endorse the move from Weak- to Strong-LOT. Remarkably, though, in a "historical footnote," Fodor and Pylyshyn (1988, 49) raise the possibility that pictures are structured vehicles for thought:

    > Connectionists are Associationists, but not every Associationist holds that representations must be unstructured. Hume didn't, for example. Hume thought that mental representations are rather like pictures, and pictures typically have a compositional semantics: the parts of a picture of a horse are generally pictures of horse parts.

    It's not clear that Fodor and Pylyshyn endorse Hume's view, but they don't explicitly reject it. And other philosophers, such as Sober (1976) and Westerhoff (2005), have argued that pictures do have a compositional semantics.
16. By 'isomorphism' I mean a structure-preserving function from elements in one domain onto elements in another. However, this function is generally understood to be selective, and may abstract away from considerable detail in each domain. I also mean to allow for a certain degree of imprecision, as when an ordinary road map reflects physical distances to within some acceptable margin of error. Such imprecisions are a ubiquitous feature of representations across the board: many, perhaps most, of our linguistic utterances are strictly speaking false but pragmatically acceptable. With respect to pictures in particular, I claim only that they depend heavily on isomorphism in respect of visual appearance, not that isomorphism suffices for depiction. As I mentioned in §1, any representational system only fixes content relative to a system of use. At a minimum, some further conditions must be added to explain the fact that depiction, like representation generally, is asymmetrical, while isomorphism is symmetric (cf. Goodman 1968, 1970).
17. Pictures may well implicitly represent non-visual features, like *being Napoleon*, or *having been abandoned by one's lover*, when they are embedded in an appropriate system of use.

18. Onomatopoeic expressions like 'buzz' are an obvious, but limited, exception.
19. Thus, in natural languages, sentences' surface structure often departs markedly from what linguists think of as their 'real', logical form, as when question-formation inverts the order of subject and verb. Pragmatic principles do exploit isomorphisms between e.g. the order of presentation and the order of described events, but these principles are cancelable, and not part of syntax or semantics.
20. It's a matter for substantive linguistic analysis exactly what the combinatorial principle(s) of natural language(s) may be: functional application, predicate modification and abstraction, conjunction, or something else entirely. Thanks to Robert May and Michael Glanzberg for pressing this point.
21. Peacocke (1997, 21) also claims that thoughts share a structure with sentences at the level of reference.
22. There are interesting exceptions: subway maps drop information about relative distances and just employ a topological equivalence, or homeomorphism, between vehicle and content. Mercator maps systematically distort spatial structure to map a sphere onto a flat plane. However, just as cubism systematically distorts but doesn't abandon the basic principles of pictorial systems, these maps retain the basic principle of spatial isomorphism.
23. Cf. also Larkin and Simon (1987). Shin (1994, 163) uses this phenomenon of automatically displaying "conjunctive information" to distinguish 'iconic' from 'symbolic' representational systems. Shin (2003, 52) extends the point to argue that because reality itself is accumulative, any system that needs to represent conjunction explicitly must be symbolic rather than iconic.
24. If maps are holistic, does a thinker who employs maps only has one (very large, complex) belief? I don't think so, and not just because a single thinker might employ multiple maps. Where a sentential model for thought individuates beliefs in terms of distinct sentences written in a 'belief box', a cartographic model should individuate beliefs in terms of minimal syntactic differences between maps. So, for instance, the thinker who endorses Figure 4 believes *that* Bob is at the grocery store (as well as that the grocery store is at 10th and South, and that Bob is less than a block from home . . .), because there is a minimally differing map in which the icon for Bob is located elsewhere.
25. Thanks to John Collins for this suggestion.
26. Ted is not at the park because the icon for Ted, 'T', is written in white, which signifies negation. Alice is either at the café or at the bar because the icon for Alice, 'A', is written twice, both times enclosed with dashed lines, which signifies disjunction. Bob was at the grocery store because the icon for Bob, 'B'', is written in italics, which signifies past tense. No one else is at the grocery store, and no one is at home, because the background for those locations are white, which signifies exclusion. And the thinker is on Lombard west of 11th because the egocentric icon, '[me!]', is located on the line for Lombard to the left of the line for 11th, with the standard orientation north = up implicitly assumed.
27. A cartographic system can capture two other aspects of intentionality: first, it can employ distinct co-referring icons—although the holistic, accumulative nature of maps makes it more likely than in a sentential system that the thinker will wonder why e.g. her 'Clark Kent' and 'Superman' icons are always co-located. Second, it can employ icons for non-existent entities, such as a 'Snuffleupagus' icon. Thanks to Adam Sennet for discussion.

28. This contrasts sharply with diagrammatic systems, which are especially well-suited to represent and relate quantificational information in a visually perspicuous way.
29. We might introduce a way to represent objects and properties as being in vague locations: say, with a blurry gradient of color, or a bubble with the phrase 'Here be *F*s' written over it. But unless such a gradient or bubble encompasses the entire map, it still attributes some location to the object/property.
30. An italic '*B*' counts as syntactically complex even though it lacks physically distinguishable parts, because it has formally distinguishable features with distinct semantic significances.
31. In keeping with the previous restriction, one would need some non-spatial means to mark the different modes of combination for *the boy bit the dog* and *the dog bit the boy*.
32. Indeed, Braddon-Mitchell and Jackson (1996, 172) argue that maps must be potentially expressively complete:

> We can be certain that something map-like can serve to represent any empirical fact about our world. The world itself is map-like: it is a vast array in space-time, rather than a two-dimensional configuration on paper, but that difference is inessential to its map-like status. And of course, the world itself makes true each and every fact about our world; it is a *perfect* representation of itself.

I don't think this is the best way to make a case for maps. As I argued in §3, even if we were to employ the entire world as a representation of itself, representing universal quantification requires an additional closure condition to the effect that the world is all that is the case. Further, many empirical facts depend upon counterfactual relations, which even a "vast array in space-time" as large as the world itself doesn't suffice to represent. Finally, in important respects the entire world is *less* expressively powerful than an ordinary Rand-McNally road map, because it doesn't have symbolic icons like 'Philadelphia' written on it.

33. Pylyshyn (e.g. 2002) has repeatedly pressed this worry against those who claim that some mental processing involves mental imagery.
34. Thanks to Bernard Molyeaux for pressing this worry.
35. Cf. Kass (1997, 107): "[T]he now widely held opinion is that the topographic features of cortical and subcortical maps are not incidental, but essential to much brain function." Likewise, Diamond et al (1999, 64) write: "The presence of 'maps' in sensory cortex is a hallmark of the mammalian nervous system... Topographical cortical representations of sensory events... appear to constitute a true structural framework for information processing and plasticity." Among many other sources, see Hubel and Wiesel 1962, Tootell et al 1988, and Van Essen et al 2001.
36. Janata et al (2002).
37. E.g. Fodor and Pylyshyn (1988, 28): "We are *not* claiming that you can't reconcile a Connectionist architecture with an adequate theory of mental representation (specifically with a combinatorial syntax and semantics for mental representations). On the contrary, of course you can: All that's required is that you use your network to implement a Turing machine, and specify a combinatorial structure

for its computational language. What it appears that you can't do, however, is have both a combinatorial representational system and a Connectionist architecture *at the cognitive level.*"

38. See Knudsen et al (1987) for discussion of computational maps, and Gallistel (1990, ch. 5) for a computational model of how a high-level cognitive map might be constructed.

References

Anderson, John R. (1978): "Arguments Concerning Representations for Mental Imagery," *Psychological Review* 85: 4, 249–277.

Barwise, John, and John Etchemendy (1990): "Visual Information and Valid Reasoning," in *Visualization in Teaching and Learning Mathematics*, ed. W. Zimmermann and S. Cunninghamdr (Washington, DC: Mathematical Association of America).

Bermudez, José Luis (2003): *Thinking Without Words* (Oxford: Oxford University Press).

Boesch, Christopher and Hedwige Boesch (1984): "Mental Maps In Wild Chimpanzees: An Analysis Of Hammer Transports For Nut Cracking," *Primates* 25, 160–170.

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

Camp, Elisabeth (forthcoming): "Putting Thoughts to Work: Concepts, Systematicity, and Stimulus-Independence," *Philosophy and Phenomenological Research.*

———(2006): "Metaphor and That Certain 'Je Ne Sais Quoi'," *Philosophical Studies* 129:1, 1–25.

———(2004): "The Generality Constraint, Nonsense, and Categorial Restrictions," *Philosophical Quarterly* 54:215, 209–231.

Carruthers, Peter (2002): "The Cognitive Functions Of Language," *Behavioral and Brain Sciences* 25, 657–726.

Cheney, Dorothy and Robert Seyfarth (2007): *Baboon Metaphysics: The Evolution of a Social Mind* (Chicago: University of Chicago Press).

Cosmides, Leda and John Tooby (1992): "Cognitive Adaptations for Social Exchange," in *The Adapted Mind*, ed. J. Barkow, L. Cosmides and J. Tooby (Oxford: Oxford University Press).

Crane, Tim (1992): "The Nonconceptual Content of Experience," in Tim Crane (ed.), *The Contents of Experience* (Cambridge: Cambridge University Press), 136–157.

Diamond, Mathew, Rasmus Petersen, and Justin Harris (1999): "Learning through Maps: Functional Significance of Topographic Organization in Primary Sensory Cortex," *Journal of Neurobiology* 41, 64–68.

Davidson, Donald (1978/1980): "Intending," in *Essays on Actions and Events* (Oxford: Oxford University Press).

———(1975): "Thought and Talk," in *Mind and Language*, ed. S. Guttenplan (Oxford: Oxford University Press), 7–23.

Davies, Martin (1991): "Concepts, Connectionism, And The Language Of Thought," in *Philosophy and Connectionist Theory*, ed. W. Ramsey, S. Stich and D. Rumelhart, (Hillsdale NJ: Lawrence Erlbaum).

Devitt, Michael (2005): *Ignorance of Language* (Oxford: Oxford University Press).

Dretske, Fred (1981): *Knowledge and the Flow of Information* (Cambridge, MA: MIT Press).

Dummett, Michael (1993): *The Origins of Analytical Philosophy* (London: Duckworth).

———(1989): "Language and Communication," in *Reflections on Chomsky*, ed. A. George (Oxford: Oxford University Press).

Evans, Gareth (1982): *The Varieties of Reference*, ed. J. McDowell (Oxford: Clarendon Press).

Evans, Jonathan St. B. T. (1982): *The Psychology of Deductive Reasoning* (London: Routledge and Kegan Paul).

Fodor, Jerry (1987): "Why There Still Has to Be a Language of Thought," in *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge, MA: MIT Press), 135–154.

———(1975): *The Language of Thought* (Hassocks, Sussex: Harvester Press).

———and Brian McLaughlin (1990): "Connectionism and the Problem of Systematicity," *Cognition* 35, 183–204.

———and Zenon Pylyshyn (1988): "Connectionism and the Cognitive Architecture of Mind," *Cognition* 28, 3–71.

Gallistel, C. Randy (1998): "Brains as symbol processors: the case of insect navigation," in *An Invitation to Cognitive Science, vol. 4: Conceptual and Methodological Foundations* 2nd ed., ed. S. Sternberg and D. Scarborough (Cambridge, MA: MIT Press), 1–51.

———(1990): *The Organization of Learning* (Cambridge, MA: MIT Press).

———(1989): "Animal Cognition: The Representation of Space, Time and Number," *Psychology Annual Reviews* 40, 155–189.

Goodman, Nelson (1970): "Seven Strictures on Similarity," in *Experience and Theory*, ed. L. Foster and J.W. Swanson (Amherst: University of Massachusetts Press), 19–29.

———(1968): *Languages of Art: An Approach to a Theory of Symbols* (Indianapolis: Bobbs-Merrill).

Gould, James L. (1986): "The Locale Map Of Honey Bees: Do Insects Have Cognitive Maps?" *Science* 232, 861–863.

Grice, Paul (1957): "Meaning," *The Philosophical Review* 66:3, 377–388.

Harman, Gilbert (1986): *Change in View: Principles of Reasoning* (Oxford: Bradford Books).

Hubel, David and Torsten Wiesel (1962): "Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex," *Journal of Physiology* 160, 106–154.

Janata, Petr, Jeffrey L. Birk, John D. Van Horn, Marc Leman, Barbara Tillmann, Jamshed J. Bharucha (2002): "The Cortical Topography of Tonal Structures Underlying Western Music," *Science* 298:5601, 2167–2170.

Johnson, Kent (2004): "On the Systematicity of Language and Thought," *The Journal of Philosophy* 101:3, 111–139.

Kaas, John (1997): "Topographic Maps Are Fundamental To Sensory Processing," *Brain Research Bulletin* 44, 107–112.

Knudsen, Eric, Sascha duLac, and Steven Esterly (1987): "Computational Maps In The Brain," *Annual Review of Neuroscience* 10, 41–65.

Larkin, Jill and Herbert Simon (1987): "Why a Diagram is (Sometimes) Worth 10,000 Words," *Cognitive Science* 11:1, 65–99.

Lemon, Oliver and Ian Pratt (1998): "On The Insufficiency Of Linear Diagrams For Syllogisms," *Notre Dame Journal of Formal Logic* 39:4, 573–580.

Marr, David (1982): *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (San Francisco: WH Freeman).

McDowell, John (1994): *Mind and World* (Cambridge, MA: Harvard University Press).

Millikan, Ruth Garrett (1998): "A Common Structure For Concepts Of Individuals, Stuffs, And Real Kinds: More Mama, More Milk, And More Mouse," *Behavioral And Brain Sciences* 21, 55–100.

Origgi, Gloria and Sperber, Dan (2000): "Evolution, Communication and the Proper Function of Language," in *Evolution and the Human Mind*, ed. P. Carruthers and A. Chamberlain (Cambridge: Cambridge University Press).

Peacocke, Christopher (1997): "Concepts without Words," in *Language, Thought and Logic: Essays in Honour of Michael Dummett*, ed. R. Heck (Oxford: Oxford University Press).

Price, Huw (1990): "Why 'Not'?" *Mind*, 99:394, 221–238.

Pylyshyn, Zenon (2003): *Seeing and Visualizing: It's Not What You Think* (Cambridge MA: MIT Press).

———(2002): "Mental Imagery: In Search of a Theory," *Behavioral And Brain Sciences* 25, 157–238.

———(1973): "What the Mind's Eye Tells the Mind's Brain: A Critique of Mental Imagery," *Psychological Bulletin* 80, 1–24.

Quine, Willard Van Orman (1960): *Word and Object* (Cambridge: Cambridge University Press).

Rescorla, Michael (2005): "Predication and Cartographic Representation," American Philosophical Association, Pacific Division, San Francisco, CA.

Searle, John (1969): *Speech Acts: An Essay in the Philosophy of Language* (Cambridge: Cambridge University Press).

Sellars, Wilfrid (1956/1997): *Empiricism and the Philosophy of Mind* (Cambridge, MA: Harvard University Press).

Shin, Sun-Joo (2003): *The Iconic Logic of Peirce's Graphs* (Cambridge MA: MIT Press).

———(1994): *The Logical Status of Diagrams* (Cambridge: Cambridge University Press).

Shimojima, Atsushi (1996): "Operational Constraints in Diagrammatic Reasoning," in *Logical Reasoning with Diagrams*, ed. J. Barwise and G. Allwein (Oxford: Oxford University Press).

Sober, Eliot (1976): "Mental Representations," *Synthese* 33, 101–148.

Strawson, Peter F. (1963): *Individuals: An Essay in Descriptive Metaphysics* (Garden City NJ: Anchor Books).

Tootell, R.B., S.L. Hamilton, M.S. Silverman and E. Switkes (1988). "Functional Anatomy Of Macaque Striate Cortex. II. Retinotopic Organization," *Journal of Neuroscience* 8:5, 1531–68.

Travis, Charles (1994): "On Constraints of Generality," *Proceedings of the Aristotelian Society*, new series 44, 165–88.

Tufte, Edward (1990): *Envisioning Information* (New York: Graphics Press).

Van Essen, D.C., J.W. Lewis, H.A. Drury, N. Hadjikhani, R.B. Tootell, M. Bakircioglu, and M.I. Miller (2001): "Mapping Visual Cortex In Monkeys And Humans Using Surface-Based Atlases," *Vision Research* 41, 1359–1378.

Wason, Peter and Philip Johnson-Laird (1972): *Psychology of Reasoning: Structure and Content* (Cambridge, MA: Harvard University Press).

Westerhoff, Jan (2005): "Logical Relations between Pictures," *Journal of Philosophy* 102:12, 603–623.