

UC Irvine

UC Irvine Previously Published Works

Title

Maps, languages, and manguages: Rival cognitive architectures?

Permalink

<https://escholarship.org/uc/item/2j56q5xd>

Journal

Philosophical Psychology, 28(6)

ISSN

0951-5089

Author

Johnson, Kent

Publication Date

2015-08-18

DOI

10.1080/09515089.2014.893814

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, availalbe at

<https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Maps, languages, and manguages: Rival cognitive architectures?

Kent Johnson

Provided we agree about the thing, it is needless to dispute about the terms.
—David Hume, *A treatise of human nature*, Book 1, section VII

Map-like representations are frequently invoked as an alternative type of representational vehicle to a language of thought. This view presupposes that map-systems and languages form legitimate natural kinds of cognitive representational systems. I argue that they do not, because the collections of features that might be taken as characteristic of maps or languages do not themselves provide scientifically useful information above and beyond what the individual features provide. To bring this out, I sketch several allegedly distinctive features of maps, and show how they can easily be grafted onto a simple logical language, resulting in a hybrid “manguage.” The ease with which these linguistic and map-like properties can be integrated into a single representational system raises the question of what work broader categories like language and map are doing. I maintain that, as an empirical matter of fact, they serve no particular purpose in cognitive science, and thus do not characterize distinct kinds of cognitive architectures.

Keywords: Cognitive Architecture; Language; Map

1. Introduction

An enduring issue in the philosophy of mind, at least since the publication of *The language of thought* (Fodor, 1975) concerns the vehicle of thought. Are thoughts sentential entities, or does cognitive activity occur in some nonlinguistic format? Over the years, numerous alternatives to the language of thought view (LoT) have been developed. One possibility frequently broached is that cognitive activity might instead employ map-like representations—in short, that thinking with maps is genuinely

Kent Johnson is Associate Professor at the University of California, Irvine.

Correspondence to: Kent Johnson, Department of Logic and Philosophy of Science, 3151 Social Science Plaza A, University of California, Irvine, CA 92697-5100, USA. Email: johnsonk@uci.edu

different from thinking in a language (e.g., Anderson, 1978; Armstrong, 1973; Bogen & Woodward, 1992, pp. 598–599; Braddon-Mitchell & Jackson, 1996, pp. 168–175; Cummins, 1996; Devitt, 2006, pp. 11, 146, & 268; Goodman, 1968; Nersessian, 2004, pp. 141–142; Rescorla, 2009, p. 176; Sober, 1976; Thagard, 2012, chapter 9; Westerhoff, 2005, p. 603).

My own view contrasts starkly with this dominant trend, primarily by rejecting a fundamental presupposition of the discussion. Distinguishing maps and languages is useful in a “loose and popular sense.” However, such talk has no analogue in a “strict and philosophical” context. Instead, I believe that in terms of mental architecture, such general categories as “language” and “map” play no significant explanatory role. (My focus here is on thought; I make no such claims for other areas of cognitive science, such as linguistics.) The various language-like and map-like properties that do serve some purpose themselves form a rather diverse lot. A great many of these properties can be consistently combined to form various sorts of representational systems, some of which are better than others for various distinct purposes. Some of these representational systems might feel more language-like to you, others might feel more map-based to you, still others might seem to be a bit of both, and so on. But once we’ve specified the particular properties of these representational systems, we’ve characterized matters fully, and there isn’t any theoretical work left to be done by declaring some cluster of potential representational systems to be linguistic, and others to be map-like, or by insisting that these latter categories not overlap. To the extent that we have characterized the details of a given representational system, we have already characterized whatever features might determine its membership in an as-yet-to-be-specified more general category, such as language. So unless it is shown that something turns on whether not a given representational system falls into such a broader category, *above and beyond the properties we’ve already explicitly laid out*, it would seem that the more exact description of the particular system is all we need, and the issue of category membership—i.e., whether the thinking is occurring in maps or in language—adds nothing further.

This attitude is not new. Like ‘language’, the term ‘innate’ is used in a wide variety of ways. In response, a prominent biologist offered the advice: “say what you mean (even if it uses a bit more space) rather than unintentionally confuse your readers by employing a word such as innate that carries so many different connotations” (Bateson, 1991, p. 22; quoted in Griffiths, 2009). In many ways, this reflects the present problem: it’s not that there’s no way to define ‘languages’, ‘maps’, etc.; the problem, rather, is that there’s far *too many* reasonable ways to do so, none of them are particularly privileged, and they yield radically different answers to the question whether a map-based cognitive architecture differs from a linguistic one.

That’s the view on offer. My case for it is as follows. Section 2 provides some background; section 3 presents several features of maps that purportedly distinguish them from linguistic expressions. Sections 4–7 shows how to build these features into a representational system of first-order logic, with the resulting system, a “manguage,” being both language-like and map-like. Importantly, there’s nothing special *or empirically implausible* about how it was put together. There are any number of ways one can cobble together various representational properties, be they language-like,

map-like, or otherwise. These possibilities are of considerable interest, but they also display how little of importance is left to the determination of the broader categories of “linguistic” or “map-like” representations. Section 8 discusses this and related points. Section 9 concludes the paper.

2. Background: Carving Out a Relevant Issue

This section sketches some general details and background assumptions of the present project. Much of this will be familiar, so my remarks will be brief; more details can be found in the references.

Generally speaking, I follow the dominant trend of adopting a (psycho-)functional perspective regarding the mind and the sense in which thinking might occur in maps, sentences, etc. (Kosslyn, 1994 is a notable exception here). Thus, my interest is in the computational character of the representations, rather than with the physical or algorithmic details of their realization; e.g., Anderson (1978), Chomsky (1965, 2000, 2002), Chomsky and McGilvray (2012), Fodor (1975, 1987, 1998, 2008), Marr (1982).

Importantly, this familiar view distinguishes a couple of psychologically relevant notions of “processing.” One notion of processing corresponds to Marr’s (1982, chapter 1) “algorithmic level” of description. This level is concerned with the (idealized) characterization of how various patterns of activation occur in the realization of some bit of cognitive activity. (These patterns are themselves approximately physically realized by the transfer of charged sodium ions, etc. in the brain.) This type of processing involves the workings of sub-representational units, and are not relevant here. However, a more abstract level of “cognitive processing” is important. A computational characterization of a creature’s psychological makeup describes how various combinations of inputs from perception, other representations, etc. result (by whatever means) in outputs of other representations, etc. That is, it describes how cognitive units are transformed during mental activity, such as reasoning, planning, evaluating, etc. Cognitive processing is a key component of the input/output structure of this type of functional-cum-computational description.¹ I take this view of cognitive processing to be relatively uncontroversial. For example, Anderson writes in his classic paper that

A representation without any process assumptions is not a theory. By making different process assumptions, it is possible to have quite different theories with the same propositional representation. We cannot test representations but only representation-process pairs. It is not an argument against a representation-process theory that there happens to be a different theory with the same representations but different processes that makes the opposite prediction. (1978, p. 257; see also pp. 250 & 263)

Virtually everyone agrees, at least in principle, that cognitive processing assumptions are crucial to computational-level characterizations of psychological phenomena; e.g., Fodor and Pylyshyn (1988). However, the importance of this point should not be underestimated. For example, Waskan (2003, 2006) shows how some computerized scale models of simple scenarios behave like mental representations of the scenes. But of course, the models don’t do this by themselves (nor do they solve the frame

problem, as Waskan suggests). Rather, they are interpretable, and exhibit any behavior at all, only when embedded in their strong processing programs. That is, the behavior in question is the result of the models, *along with* the strong knowledge of basic physics encoded in their host “off the shelf” commercial software (Waskan, 2003, p. 270; 2006, p. 183). Thus, it is only in the context of the processing software (which, for economic reasons, incorporates the types of physical knowledge its users are likely to exploit) that one can claim that “these representations implicitly contain all the information needed to predict the consequences of countless alterations to their represented systems” (Waskan, 2006, p. 193). Section 7 below offers further discussion of this important assumption.²

The LoT hypothesis maintains that, at an appropriate level of description, cognitive activity occurs by constructing and manipulating sets of sentences. The authors mentioned above in section 1 hold that thinking with maps is something else altogether. Whether they are right is unclear, since there’s no consensus about what constitutes a “sentence” or a “language,” not even in logic. For example, in automata theory, a language is simply any (non-empty) subset of the set of all (non-empty) finite strings of elements from some given finite set (e.g., Ebbinghaus & Flum, 1999, chapter 6). In model theory, the sets of primitive expressions and the lengths of formulas are often allowed to be infinite; e.g., Hodges (1993, chapters 1–2). (Notice, incidentally, that the automata theory definition does not require that a language enjoy any such properties as compositionality or systematicity, but does require that sentences be finite; the model theory definition does just the opposite. Thus, even these sacred properties of languages are not always preserved.)

These are just two examples; many more could be supplied. For example, Berwick and Chomsky write that the term ‘language’

has engendered much confusion. Sometimes it is used to refer to human language; sometimes it is used to refer to any symbolic system or mode of communication or representation, as when one speaks of the language of the bees, or programming languages, or the language of the stars, and so on. (2011, p. 19)

For similar sentiments, see Cela-Conde and Marty (1998, p. 24), Chomsky and Stemmer (1999, p. 397), and Collins (2007, p. 651). (This inspecificity of ‘language’ is a major reason why Chomsky and others have adverted to very particular and specialized notions of ‘language’—more precisely *i-language*; see Chomsky, 1986, 2000—when discussing generative linguistics and its subject matter.)

Regardless of what is to count as a language, the general project of specifying the broad structural or “architectural” features of a representational system centrally involves identifying, at some useful level of description, the general means by which some interesting range of cognitive activity occurs. What properties such a system has, and what desiderata it should satisfy, have been subjects of enormous controversy; see Churchland (2012), Elman et al. (1997), Fodor (1975, 2008), Fodor and Pylyshyn (1988), Jackendoff (2002), and van Gelder (1995). A corollary of this controversy is that any such vehicle of thought is not open to direct introspection; gut feelings are not relevant here (for example, Sober writes, “I suggest that it is possible that the *form* of

representations used in cognition is opaque to introspection, even though the contents of particular cognitions may be somewhat transparent,” 1976, p. 111). Similarly, the phenomenological sensation of visualizing some region of space “in the mind’s eye” does not necessarily reflect whatever mechanisms ultimately underlie the cognitive activity in question; that remains a challenging empirical matter. As Pylyshyn puts it, “we are not entitled to take the tempting road of assuming that the content of our experience reflects in any direct way the nature of our cognitive information processing activity” (2002, p. 180; see also Gattis, 2001; Kosslyn, 1980; Thomas, 2011; Tye, 2000).

Lacking an antecedently given and commonly accepted notion of a language (and similarly for maps), the broader task of assessing whether thinking in maps differs from thinking in language is a much more abstract affair. The import and interest of this task relies on the import and interest of what counts as a linguistic entity, and what counts as a map. Unfortunately, some plausible distinctions simply dissolve the issue before it begins. For example, assuming the phonological properties of natural languages are part of their structure, few would oppose the claim that thinking in map-systems is distinct from thinking in English. Indeed, some of the strongest proponents of the view that thought *must* occur in a language also hold that thought *cannot* occur in a natural language (e.g., Fodor, 2001, 2008, chapter 3; Pylyshyn, 2003; although see Berwick & Chomsky, 2011 and Chomsky & Stemmer, 1999 for some additional interesting speculations about natural language and thought). In this sense, it’s trivially *true* that thinking with maps is distinct from thinking with “language.” In the other direction, we might take a cue from van Gelder (1990, p. 360) and hold that a representation counts as a sentence if its constituents are tokened whenever it is tokened (see also Carruthers, 2006, section 1.6.2, and Fodor & Pylyshyn, 1988; van Gelder’s “concatenative compositionality” is more sophisticated than this in ways that will not matter here).³ Virtually everyone who holds that thinking with maps is distinct from thinking with language agrees that maps have constituents that are tokened whenever they are tokened. So with these characterizations, it’s now trivially *false* that thinking with maps is distinct from thinking with “language.” In short, to get a worthwhile discussion of the issue off the ground, we need appropriate notions of maps, languages, and the like. The question is whether cognitive science provides them.

3. Some Noteworthy Features of Maps

This section describes four features of maps that have been used to distinguish them from linguistic objects. These will suffice to structure the ensuing discussion; later sections consider some other ways they have been distinguished. Importantly, neither these four nor any of the others discussed below are intended to constitute a conceptual analysis of maps. Perhaps someone can find a feature that decisively separates systems of maps from languages in a way that matters. But I don’t know of any such feature or the case to be made for it. So at the very least, this section puts the burden of explanation on those who hold that there must be some important difference.

Sober (1976) suggests that pictorial or map-like systems differ from linguistic ones because only the former are continuous, in the sense of being able to supply a continuum of representations of length, color, area, etc. Indeed, he considers this to be the key distinguishing feature that renders pictorial representations as non-linguistic entities (1976, pp. 138–146).⁴

Camp (2007) similarly distinguishes maps from language by holding that only maps combine elements—e.g., symbols for schools, roads, lakes, etc.—according to a principle of “spatial isomorphism.” The basic idea is this. Your typical map is a two-dimensional object containing various symbols at various distances from one another, and these distances in the map necessarily represent distances in the world. The fact that the symbol for the hospital is nowhere near the symbol for the school on the map indicates that the hospital itself is not next to the school in real life. Similarly, the fact that the road goes from points A to B and then to C indicates that it doesn’t go from A to C and then to B.

The use of distance and proximity in maps to represent distance and proximity in real life contrasts with the use of distance and proximity in language, according to Camp. Unlike maps, linguistic systems

abandon any sort of direct physical isomorphism between vehicle and content The syntactic principles combining those constituents . . . clearly abandon any appeal to physical isomorphism. 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. 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. (Camp, 2007, pp. 156–157)

Finally, maps are often held to be “holistic” representations, capturing all of the relative spatial locations of all of the entities represented. For instance, Braddon-Mitchell and Jackson write that “each part of a map contributes to the representational content of the whole map, in the sense that had that part of the map been different, the representational content of the whole would have been different” (1996, p. 171; see also Camp, 2007, pp. 161 & 167; actually, these issues are a little more complicated; see Rescorla, 2009, p. 195). That is, if a map changes because the church-symbol is now north of the school-symbol, this also changes the relationships the map represents the church as bearing to the river, the forest, etc.

Thus, we can identify the following four features of systems of maps:

- (1) i. *Continuity*: Systems of maps are capable of varying continuously.
- ii. *Spatiality*: Maps have a semantically interpretable, (two-dimensional) spatial aspect that relates the represented elements.
- iii. *Interpretation*: This spatial aspect is necessarily interpreted as spatial distance (in a way that, together with the previous condition, results in the spatial isomorphism mentioned above).
- iv. *Holism*: The spatial aspect is “holistic,” in that all distances and angles between the represented elements are presented.

The authors mentioned above (as well as numerous others; see below) argue that the properties in (1) distinguish map-like systems from linguistic ones. In the next few sections (sections 4–7), I will develop a representational system (a “manguage”) made up of these four properties and a simple logical language. The results are summarized in (5)–(8) in section 5. As we’ll see, given what we currently know, it wouldn’t be at all surprising to discover a creature that employed something like it as part of its cognitive architecture. This will show that the possibility sketched below is not too far-fetched, lurking only in some distant bit of logical space. Similarly, it further undermines the importance of a distinction between thinking with a language versus a system of map-like representations.

I (re-)emphasize that (1) is not meant to be a conceptual analysis of any sort. The point is only that these are some of the more prominent features that the literature has claimed distinguish map-thinking from language-thinking. The following sections show that they don’t do a good job at this, since it’s easy to combine them with language-like features to construct representational systems that, as an empirical matter of fact, are perfectly viable.

4. Continuity, Maps, and Languages

Imagine that we are studying the cognitive capacities of some creature—human or otherwise—and we’ve learned that some functional properties of the creature’s thoughts correspond to (1). Our task is to further develop our theory of its representational system. If the authors cited in section 1 are correct, we should expect any such resulting system to support an empirically important rift between the two types of thought they distinguish. In particular, the resulting system should not be language-like. I take the features in (1) in turn.

Despite its common appeal, issues of continuity are a red herring. However, two matters should be distinguished. On the one hand, there is the empirical phenomenon of the creature’s own representations of spatial relations, etc. On the other hand, there are the scientific community’s representations of this empirical phenomenon, often via formal mathematical models. In particular, inherently discrete phenomena are often modeled with real numbers, so as to employ the powerful tools of differential equations, etc. Thus, theorists might, as a convenient but false idealization, characterize the creature’s representational abilities as continuous, even though e.g., it’s representations fail to discriminate distinctions beyond certain thresholds, thus constituting a discrete system of representations.⁵ There are some subtle issues here, but at present the relevant issue is not whether researchers use continuous mathematical structures to characterize the empirical phenomena under study. Rather, the issue is whether the empirical phenomenon itself—i.e., the creature’s representational system—varies continuously in the relevant ways. Moreover, this matter concerns whether the creature’s representations make relevant discriminations, not the further issue of whether the creature subsequently reacts in some detectable discriminatory fashion.⁶ In what follows, I will assume that the creature’s *representations* do vary continuously, perhaps because some significant component,

like the wavelength or total energy of (some part of) the internal representation, varies continuously.

Even if the creature's mental representations vary continuously, they may still take the form of sentences in a logical language. As a first pass (generalized below) we might, e.g., consider the creature as having the capacity to realize, perhaps by constructing on the fly, any of the continuum-many two-place relation-symbols R_{ax} , which have the interpretation x is a far away from y , for the distance magnitude a (presumably non-numerically represented by the creature). That is, although the creature only realizes finitely many cognitive events, its cognitive capacity (or "competence" in the technical Chomskyan sense) involves the means for assembling any of these relations "online," as partly a result of, and/or as part of, the process of perceiving and engaging in spatial cognition. Or, what may amount to the same thing, we can think of the creature as having a cognitive spatial predicate with a demonstrative-like parameter that indicates that x and y are *this* far apart (see Peacocke, 1986 for some similar sentiments). In either case, the creature sees x as *yea* far away from y , which is part of the cause of it forming the perceptual representation $R_{yea}xy$.⁷

In short, there's no reason to suppose map-like capacities can vary with a level of resolution not attainable by language-like ones. (Indeed, utilizing languages with a continuum of first-order nonlogical symbols has an impressive history. They were the basis behind Robinson's (1961) proof of the existence of models of nonstandard analysis.)

5. The Spatial Aspect of Maps

Condition (1.ii) says that map elements are spatially extended. Logic standardly designates relations with sentence letters. However, there is no reason, aside from practical convenience, why a logical language might not instead use two-dimensional regions (i.e., distances) of varying length and angularity relative to some fixed direction, all of which are understood to vary in the same two dimensions. Consider, for example, some formulas of the form $R_{ab}xy$, interpreted as " y is a far away from x at an angle b (relative to some designated direction, such as left, due north, etc.)." Call this the *oriented distance* from x to y . We can now replace the symbol R_{ab} with a proportional spatial region. Thus, the formulation $R_{ab}xy$ is now shorthand for the official syntax:

$$(2) \quad \begin{array}{c} y \\ x \end{array}$$

In (2), the particular syntactic relation symbol is determined by the oriented distance between x and y (and whatever scaling relation exists between the representation and represented scene). Notice, incidentally, that this representation does not require any coordinate system; moreover, if direction is unimportant, the designated direction is unnecessary. (Such nondirectionality could be implemented with three-place relations

Rxyz that give the angle formed by xyz, and length of the sides. The rigidity of the representation—i.e., invariance under translations, rotations, and reflections—is insured by there being a distinct formula for all three distinct elements of the map; see below.)

The other spatially extended elements of maps, such as roads, regions indicating forests, buildings, etc., are handled similarly, taking constant symbols to be themselves spatially extended. For example, the statement, which in a different notation might be written Rxy, that we will later interpret as the road passing through that part of some particular region, might be the relational formula:



This satisfies the spatiality condition of (1): the representations possess a (two-dimensional) spatial aspect that relates various objects and their parts.

There's nothing inherently two-dimensional about the spatial representation just offered. The "maps" in question can just as easily be three-dimensional scale models (Waskan, 2003, 2006), or even higher-dimensional representations. Although I will continue to use standard two-dimensional maps in the discussion, everything that is said about them applies to more complex structures with an n-dimensional (not necessarily Euclidean) geometry.

At this point, one might naturally wonder what symbols should be used. Is a meeting of five roads at a roundabout one complex symbol, or is it six: one for each road, and one for the roundabout? This is a matter for empirical psychology; after all, it's well-known that object individuation is a very basic component of our mental lives; e.g., Carey (2009), Hatfield (2009), Xu (1997). But there is also a pleasing convergence of philosophical and psychological attitudes. From a very different perspective, Westerhoff adopts precisely this stance when he distinguishes "constituents" from "parts":

The constituents of the state of affairs are now precisely those parts which we use in conceptualizing this bit of the world; the elements which we take to correspond to the bit of language we use to identify it. But the mere parts (John's brain as opposed to John) are just any parts of that particular bit of the world we happen to be talking about, whether they take part in our conceptualization or not. In analyzing a state of affairs we therefore have to concentrate on its constituents, on the parts we actually use in thinking about this subsection of the world, rather than on any arbitrary collection of parts it happens to have. States of affairs are parts of the world which the human mind has sliced, and in investigating them we have to concentrate on the cuts the mind has made, rather than on any other fissures we may otherwise observe in the material. (2005, p. 608)

In what follows, then, I will take the individuation of the items to be spatially represented as given by the cognitive makeup of the creature under study. First, though, a word of caution: assuming that we can say what such notions as language, map, etc. amount to, there remains the further difficult (and frequently ignored) matter of saying what it would be to think "in" one such system or another. After all,

as we saw above, a theorist's representation of a creature's representation may differ in important ways from the creature's representation itself. A theorist's own theoretical characterization might involve some apparatus (which might look to you like a system of pictures, or sentences, or whatever). However, it's another matter altogether to hold that this theoretical representation captures key aspects of the creature's own representations that are not captured by a (seemingly equivalent) theoretical rival form of representation.

Consider, e.g., how a theorist might represent a creature's representation of rotation. A theorist's map-like representation might take the form of a change in the relationship between the represented object(s) and some (possibly implicit) reference axes, etc. Alternatively, the theorist's representation might take the algebraic form of an orthogonal transformation on an algebraic description of the object(s) location(s). For a great many purposes, these "map-like" and "sentence-like" representations are equivalent. The claim that the creature thinks "in" one of these representational formats and not the other must be justified in terms of properties they don't share. But given that these are "computational" level characterizations of the creature's representational capacities, it becomes extremely difficult to say what renders them plausibly non-equivalent in a theoretically interesting way. Unfortunately, everyone, bar none, who holds that thinking "in" map-systems versus languages is importantly different owes just such a story. That is, it's not enough for map-systems and languages to differ; instead, these differences must have theoretical bite at the highly abstract computational level where they do their work. This obligation appears in multiple places, too. For instance, if you adopt the very common claim that only maps are "iconic," differing from sentences "in virtue of their resemblance to what they represent" (Blumson, 2012, p. 414; Waskan, 2003), then you need to say what this resemblance amounts to at the computational level of characterization (e.g., Shepard & Metzler's 1971 famous study of mental rotation). Since I don't think that our scanty present empirical understanding supports any such distinction, I owe no such story. Indeed, a central purpose of the sketch of the manguage below is illustrate just how little the presently available empirical facts say about which properties—the allegedly "map-like" or "language-like" ones, or various admixtures—are plausibly part of the creature's representations (as opposed to the theoretician's). (I thank Jonathan Waskan for encouraging me to clarify this point.)

6. The Spatial Aspect Is Interpreted Spatially

The previous section showed how to isolate spatially extended symbols, and use them in a logical language. This section concerns their interpretation. (1.iii) says that these symbols' spatial aspects are essentially interpreted spatially. The primary task here is quite simple. However, condition (1.iii) raises some further issues for maps and languages (and manguages) which themselves merit further consideration. I start with the easy part.

Maps, it is claimed, are distinguished from other types of spatial representations, such as Venn diagrams, scatterplots, etc. in that the spatial features of the map represent spatial distance and not something else; (Camp, 2007, pp. 153–154).


A scatterplot summarizing some experimental data might contain a dot representing a particular rat, where the dot's position on the x-axis represents the amount of sucrose it was given, and the y-axis represents its performance on some task, and distances between dots are interpreted accordingly. This nonspatial interpretation disqualifies the scatterplot from being a map.

For now, we can assume that maps are in part identified by their interpretive or semantic properties, whatever precisely those may be. They are diagrams with additional semantic restrictions.⁸ That is, a map is a diagram with all spatial magnitudes, shapes, etc. interpreted as representing spatial magnitudes shapes, etc. of the same form. In logic, such restrictions are often given by specifying the class of admissible models of the system.

What's good for the goose is good for the gander. To whatever extent maps are distinguishable from the more general category of diagrams by restricting their admissible interpretations, any other representational system's admissible interpretations can be similarly restricted. Such a restriction is precisely as stipulative as the means by which maps are distinguished from diagrams (and the likes of (4) below, if such be the case). Moreover, since the present project is the construction of a kind of cognitive representational system, it is appropriate to design it to operate in some ways and not others.⁹

In one sense, the condition (1.iii) of spatially interpreting the spatial aspect of maps is mostly a technical issue admitting of a technical solution. However, in another sense, this condition is more subtle than it might initially appear. I end this section by mentioning two issues that merit further consideration.

Condition (1.iii) and the notion of a spatial isomorphism are both highly underspecified. For example, the spatial interpretation seemingly needs to be "of the right sort," because isomorphisms of spatial structure are often all too easy to obtain. For example, (4b) is a representation in Cartesian coordinates of the polar coordinates of (4a):



Properly interpreted, the two structures in (4) are isomorphic, although the distance metric in (4b) is complicated. (For example, it is not isotropic: identical Euclidean distances between two pairs of points in (4b) will not in general be identical in (4a).) However, I assume that most map-theorists would not regard (4b) as a map of (4a), despite their spatial isomorphism. Thus, (4) raises the question of just what spatial interpretations should count as admissible for a representation to count as a "map." For example, if one rules out an isomorphism between the two on the grounds that the same distance in different regions of (4b) are interpreted differently, one might thereby be forced to similarly rule out the very possibility of a map being plotted on a logarithmic scale. Worse yet, Mercator maps, perhaps *the* exemplar of maps, would no

longer qualify as maps! (Such a restriction would have other consequences too; e.g., for Hatfield's three-dimensional transformed visual space; see Hatfield, 2009, chapter 6; cf. Suppes 1977).

Unsurprisingly, I think the issue raised by (4) is a non-starter. If a creature employs representations like (4b), then we may want to know how and why it does so; e.g., perhaps (4b) is a particularly efficient way for it to store and/or process spatial information. But once these matters are sorted out, I see no *further* benefit to be had by determining or legislating whether or not the creature is thinking with maps (or with sentences, or with both, or neither). Provided we agree about the finer-grained details of the representational system, as Hume might say, it is needless to dispute about whether (4b) is a map.

The second issue stemming from (1.iii) is that it is a particularly simple way of enforcing a certain kind of interaction between semantics and syntax.¹⁰ But considerably greater complexity is possible, both theoretically and in real life. For instance, Johnson (2001) develops an extension of first-order logic, where syntax and semantics interact in various complex ways. In particular, the semantic properties of various expressions are allowed to partly determine which syntactic constructions they can enter into. Specific instances of this general phenomenon also appear in natural languages like English. For example, the subtle semantic features that result in a phrase's "thematic" structure account for the grammaticality of, say, *These vases break easily* and the ungrammaticality of **These zombies resemble easily* (i.e., it is easy to resemble these zombies); e.g., Hale and Keyser (1986, 1987). Thus, the formal possibilities and empirical realities of languages belie the hopeful thought that "in linguistic systems, this distinction between syntax and semantics is clear" (Camp, 2007, p. 155). Instead, this distinction is so unclear as to be a very active area of research.

7. A Holistic Interpretation

Condition (1.iv) says that maps have a "holistic" quality: *all* distances and angles between the represented elements are present in the representation. There are numerous ways to accommodate this feature. However, before turning to the perhaps most interesting version, it will help to first pause and say a few words about where we are.

In terms of characterizing a representational system that we might ascribe to the imagined creature under study, we have so far focused mostly on only one part of it, namely the structure of the representations themselves that the creature might generate. But, as stressed above (section 2), without a theory of how these representations are employed, there's little to be said about how well the system characterizes the creature's representational capacities. If, metaphysically speaking, the creature's representations have all the strength of first-order logic, but the processing system that uses them can only manipulate the sentences' propositional structure, then *de facto* you are dealing with propositional logic. (We noted this earlier, when observing that a creature might be able to generate continuously varying representations, but could employ them only in a more limited fashion, e.g.,

discriminating two of them only when their total deviance surpassed some threshold.) Similarly, if the processing system can handle all of the first-order structure, but also encodes a lot of information about the meaning of the nonlogical terms, etc., it may be more powerful than a purely logical deductive system. But however humans and other complex creatures process information, it is considerably more complex and less well-understood than, say, the simple deduction systems of ordinary logic.¹¹

A priori, both language-like and map-like representational systems, along with their potential associated processing mechanisms are extremely flexible, and can accommodate a very wide range of facts. Thus, in advance of further empirical details about the creature, a representational system can respect the holism constraint simply in virtue of how its representations are constructed and processed. In particular, in developing a theory of the creature's map-like cogitations, we might view the processor as always constructing formulas representing the spatial relations for all pairs of objects it represents. In such a case, the representation will be, intuitively, a complex that spatially relates the relevant objects that the creature has individuated.¹² This could be a very natural thing for a creature's mind to do—quite plausibly without the creature's conscious awareness of any such processing. For example, it might produce the representation by simultaneously developing initial estimates of all the individual spatial relations, and then improving their accuracy by balancing them off one another. Such estimates could be held in a cognitive "buffer," only becoming available to further cognitive processing after all the estimations had been finalized, at which point all of them are made available in a holistic fashion. Indeed, such checks would aid in the correction of particular errors, and might thus be evolutionarily plausible, particularly if the creature often needs to represent spatial configurations that are not present.

There are various ways a cognitive system might realize this kind of representational holism, although in advance of more empirical details, we can only speculate. One tantalizing thought is that, since by hypothesis the symbols in question (for spatial relations and the objects they relate) are themselves spatial, the system might compress the information in the holistic representation (a long string of conjunctions) by a process of "overlapping," whereby a symbol designating an object is tokened only once, and it simultaneously occurs with all the relation symbols that copies of it originally occurred in. Intuitively speaking, the system "overlaps" all the tokens of each symbol type on top of one another, where the types correspond to the distinct token objects represented. For example, if the creature represents two tokens of the same type of thing (two identical apples, say), then the token representations of the first apple will be of a distinct type from the tokens of the second apple. Such a process would of course be a form of chunking; e.g., Miller (1956).

Let's summarize sections 4–7. We began by assuming that the creature's psychological capacities individuated some objects, and formed representations of their spatial relations to one another:

$$(5) \quad R_{ab}xy \ \& \ R_{cd}xz \ \& \ R_{ef}yz$$



We then attributed distinctively spatial characteristics to the nonlogical terms, both names:

(6) R  & S  & T 

and relations:

(7)  &  & 
 

Finally, via a process of overlapping, the creature's ability enabled it to compress all this information into a single chunk:

(8) 


(8) has the structure of a simple map. As its development shows, it also has the structure of a language. Whether the creature is said to think “in a language” or “in a map” (or in both, or in neither) is of little concern once the nature of its representations and processing capacities have been articulated. Obviously, more complex such representations could be made by starting with a greater array of conjunctions in (5).¹³

I end this section with four comments. First, the process of overlapping could be a natural form of data compression, whereby a collection of formulae that the creature's processing system had evolved to treat as systematically related are collected together so as to minimize the number of symbols employed. We engage in this kind of lossless data compression all the time: e.g., we might state the ordering of some integers as ‘ $0 < 1 < 2 \dots < 8 < 9$ ’ using only 19 characters instead of the logically more correct ‘((((((($0 < 1$ & $1 < 2$) & \dots & $8 < 9$))’ which requires many more. Here too, we've overlapped the dual occurrences of eight of the symbols, using what we know about the transitivity of ‘ $<$ ’, the associativity of ‘&’, etc.

Second, as we've seen, lacking details of how representations are processed, we are in no position to conclude that “maps provide an exceptionally efficient and computationally tractable system for representing and reasoning about the world” (Camp, 2007, p. 175; see also pp. 162 & 167; see Anderson, 1978, p. 256 for diametrically opposed sentiments).¹⁴ After all, a map corresponding to the representation in (8) might be an unwieldy cognitive object that is difficult or impossible for the creature to process directly, but must first be decomposed into its non-overlapping conjuncts.¹⁵ Similarly, although the polar coordinate representation in (4) seems complicated from the perspective of our person-level overt reasoning abilities, a creature (indeed, perhaps a human one) might use it because of its exceptionally efficient and computationally tractable properties from the standpoint of its relevant processing capacities.

Third, our ignorance regarding how the processing systems interact with the representations can make it hard to correctly interpret certain kinds of behavioral evidence (e.g., Anderson, 1978). For example, the processing mechanisms considered

above may involve a consistency check that erases the representation and starts over whenever it detects even a small chance that the representation was incompletely formed. Such an evolutionary development would be unsurprising if representations are easy to make, and omissions costly for the creatures. In this case, the processing of representations (and the measurable effects thereof) might seem quite map-like to the scientist studying the creature's representational capacities. Importantly, of course, the same sort of procedures could also easily apply to even more straightforwardly "linguistic" representations of spatial phenomena.

Fourth, there's nothing empirically unusual or implausible about the representational system just constructed. Similarly, the little that is known about cognitive representation fails to undermine any number of other candidate systems of representation. Many such systems can be quite difficult to discern in terms of their measurable predictive consequences. Suppose, for instance, that some other creature at times produces genuinely map-like representations. Suppose also that it employs these representations via some form of binary comparison mechanism. Thus, its employment of these maps would occur by selecting pairs of object-representations and extracting information about the oriented distance of their representanda. In this example, the map-like representations would function in the creature's larger cognitive economy just like a large collection of conjuncts in a language, with a summary running from (8) to (5). (With a bit of imagination, its language-like properties could be further augmented.¹⁶) But here too, there doesn't seem to be any interesting issue about thinking in "maps" versus "language."

8. The Status of Manguages

What are we to make of manguage summarized in sections 4–7? There are at least three reactions one might have; I take them in turn.

The counterexample conclusion. This reaction interprets the manguage as a counterexample to the commonly held claim that thinking in maps is perforce distinct from thinking in language. One might hold that the representation in (5)–(8) has both linguistic and map-like structure, thus contradicting such claims as "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" (Camp, 2007, p. 155).

To be compelling, this view requires both that (i) we share sufficiently sharp and stable intuitions about maps and languages to agree that manguages and their ilk are counterexamples, and that (ii) these shared intuitions somehow matter. Regarding (i), I don't really know what the empirical facts are, although about equal numbers of people I've talked to have declared that the manguage sketched above is either "obviously" or "obviously not" a system of maps (as distinct from a language). This is consistent with many trends in the psychology of judgment and in experimental philosophy. Regardless, (ii) is even more suspect, particularly for establishing scientific natural kinds, psychological or otherwise. This is one of the most familiar points of all from the philosophy of science, so it needn't be belabored here.

The vagueness conclusion. This reaction interprets representational systems like languages as showing that the distinction between maps and languages is vague, with a fair amount of overlap in between these notions.¹⁷ As we've seen, a bit of thought suggests a wide variety of further representational systems that possess both linguistic and map-like properties (where these terms are used informally).

Although *prima facie* plausible, it's not obvious what this reaction amounts to. Paradigmatic vagueness cases show that there are unclear examples that lie "between" categories, where "betweenness" is realized along some unidimensional dimension of (at least approximately) continuous variation, such as the number of hairs on a person's head for baldness, the number of grains of sand for heaps, etc. To make sense of the vagueness conclusion, one would have to address the fact that there is no such ordering of representational systems from languages to maps. For example, would language A that lacked the continuity property be closer to a language than language B that could interpret the spatial aspect non-spatially? What if B interpreted the spatial aspect in the polar coordinate fashion of (4)? In what ways could continuity fail of A to bring it exactly as close to a language as B? There is simply no principled way to adjudicate these matters. Thus, talk of vagueness here seems little more than metaphorical. Worse yet, it's not enough that some such characterization of the language-to-map "spectrum" be made in just any way; rather it must be principled and relevant to the relevant part of cognitive science. This leads to the final reaction.

The pragmatic conclusion. I endorse this reaction. Like the vagueness conclusion, it is based not on the particular language of sections 4–7, but the fact that it illustrates how any number of further types of hybrid systems could be similarly constructed. Such myriad possibilities raise the question whether there is any pressing need for theories of broad classes of representational systems, such as map-systems and language. In the sense that matters for the many authors cited at the beginning of this paper, this would amount to characterizations that are sharp enough to allow us to determine whether thinking with maps really is different, in some theoretically important sense, from thinking with language. But if you know the creature under study thinks in a language as characterized above, what *more* do you learn by declaring that it thereby thinks in a language, or a map-system, or whatever? I see nothing at stake in such an issue. Of course, there might be an important distinction here; but a case for this must be made. Once the relevant features of a representational system have been articulated, there doesn't seem to be much work left to be done by subsequently categorizing them as languages or not. In some sense this isn't too surprising, since the characterization of any representational system is in terms of the finer-grained, more specific features that would themselves determine whether it falls into a broader category such as a language. But if these broader categorizations do no *additional* theoretical work, then there is little reason to regard them as reflecting genuine natural kinds of representational systems.

These considerations apply, of course, "all the way down," as would be expected. As part of the ordinary course of scientific progress, we can anticipate that many of the more fine-grained properties, e.g., systematicity, holism, etc. will themselves be similarly resolved. Representational systems that seem largely systematic, holistic, etc.

in their spatial representations may require further fine-tuning, qualifications, etc. The net result, we should expect, is a more detailed characterization of the facts, albeit one which may not sit happily with many of the well-known “loose and popular” notions that roughly but inaccurately describe the system.

For example, many artificial languages and manguages are both systematic and compositional in a very traditional sense, e.g., Fodor and Pylyshyn (1988); but they needn't be. Depending on its ecological niche and evolutionary demands, there's no a priori reason to suppose that a creature's representational capacity enables it to represent that the branch is circling around the mosquito, or that the mound is climbing up the enemy ant, although the converses may be part of its conceptual repertoire. More generally, pace Blumson (2012), systematicity is a much trickier notion for naturally occurring representational systems; see Johnson (2004). English contains all sorts of violations of our intuitive notions of systematicity (e.g., it's normal to say *Mary stowed her gear* or *John put his gear down*, but you can't switch these verbs: **Mary put her gear*; **John stowed his gear down*.) Natural languages are the best empirically occurring representations of human cognition. Barring further evidence, then, we should assume that whatever the “systematicity” of thought amounts to, like language, it will be much more complicated and unintuitive than is suggested by the standard simple examples (of the *Mary kissed Bill/Bill kissed Mary* variety). Indeed, if Johnson (2004) is correct, ‘systematicity’ itself should join ‘language’, ‘map’, etc. in the dustbin of loose and popular notions that are too general and unspecific to play the roles of genuine scientific natural kinds.

In short, lacking an appropriately motivated characterization of the key terms, it's hard to agree or even disagree that maps constitute a demonstration “in the starkest possible terms” that “there are *nonlinguistic* [emphasis added] representational vehicles composed of recurring discrete parts” (Camp, 2007, p. 175). Instead, the proper response toward the common trend of holding that thinking is maps is substantively different from thinking in language is simply to reject the terms of the discussion, at least until nature herself decides to endorse these terms.

Although I currently don't see any motivated distinction between languages and map-systems, this is a purely empirical matter. What might such motivation look like? Ultimately, justification can appear in any number of ways; here's a very common one. Suppose that while investigating various forms of representational systems of various species, it was discovered that a certain congeries of intuitively map-like properties tended to cluster together: if a type of cognitive activity (perhaps within or across species) presents some of them, it reliably presents all of them. Suppose also that this holds for some very different congeries of intuitively language-like properties. In such a case, we might legitimately identify two general kinds of mental representation, and maintain that thinking in the one really is importantly different than thinking in the other. Moreover, we would predict the non-existence, as a matter of empirical fact, of manguages.

The scenario just imagined involves a suspiciously clean pair of confluences. In real life, a more plausible scenario would be that the map-like systems tended (but didn't always) have some map-like properties, along with a few strange

other ones, and oddly tended to omit a few “obvious features”; similarly for the language-like systems. Moreover, their separation might only be a statistical trend (across species, perhaps), with many exceptions. In such a case, the legitimacy, utility, importance, etc. of these cognitive kinds would be summarily lessened by the complexities of reality. More importantly, the justification for positing these categories comes entirely from the work these kinds do. Importantly, contrary to a suggestion by Blumson (personal communication, 2013), there is no need to say what cognitive representations must be like, or cannot be like. Perhaps human mental representation may turn out to suit your feelings about what a language that is not a map-system is like, or vice-versa, or neither. But these are matters that should be based heavily on the empirical facts. At present, these facts are not strong enough to determine a computational-level cognitive architecture that has one of these forms.

9. Conclusion

The commonly held view that thinking with maps is substantively different than thinking with language rests on notions such as maps and language being genuine cognitive kinds of mental representations. But there’s presently no compelling empirical motivation that supports positing the relevance of representational systems that would support or undermine such claims.

To some extent, all this should be unsurprising, for at least two reasons. First, the space of viable theories remains quite large: the highly abstract and poorly understood subject of mental architecture provides few constraints for ruling out hypotheses. Second, in many respects, dropping these categories and any distinctions between them leaves intact all the original issues involving the compositionality, systematicity, continuity, holism, etc. of various kinds of representation. The only real change is that we recognize that when we gesture at these issues with terms like ‘language’ and ‘map’, the latter needn’t—and don’t—acquire a life of their own above and beyond the more specific issues that really do matter.

Acknowledgements

I am grateful to Jonathan Waskan and Ben Blumson for many thoughtful comments on this paper.

Notes

- [1] Talk of functional descriptions in terms of inputs and outputs is ultimately too narrow, since in many cases such descriptions crucially involve response times, patterns of development and breakdown, etc.
- [2] This off-the-shelf software is obviously implemented in a straightforward classical computational format. Thus, the sense in which such models provide a genuine alternative to a language of thought architecture is a delicate matter; see Waskan (2006, chapters 4–6) for detailed discussion.

- [3] I should stress that when suggesting that the maps-versus-language debate needs different conditions on language, I do not mean to suggest any criticisms of, e.g., van Gelder's proposal; just the opposite, in fact.
- [4] Actually, Sober is a bit ambiguous here, sometimes referring to the (unqualified) "digital character of language," other times restricting attention to English or to "linguistic systems of a certain kind," which "can be more impoverished" (1976, pp. 138–139). We will see below that it's important to get clear on this matter.
- [5] Contrariwise, a creature might employ a continuously varying type of representation in only one of two ways, depending on whether the continuous feature(s) exceed some threshold. In that case, researchers might nevertheless characterize such continuous representations as having only one of two discrete values.
- [6] Positing continuously varying representational resources has independently occurred in the literature, e.g., Gallistel and Gelman's (2000) continuous "accumulator"; see Laurence and Margolis (2005).
- [7] Such a view, including its development below, is similar to one held by Bach (1987, p. 21), although I here extend it to spatial relations between objects.
- [8] It's a little unclear whether Camp views maps as subsets of diagrams, or whether they are distinct forms of some more general category of spatial representation, call it the maps-or-diagrams. In the text, I will assume the former, but nothing will hinge on this, as we can always advert to the latter.
- [9] Similarly, it is also part of the very nature of the project to determine other aspects of the creature's representational repertoire. For example, in addition to being able to represent the oriented distance between the house and the road, it may also have expressions that enable it to identify the road as this particular road, or the road that the dog was running alongside yesterday, etc.
- [10] Or between semantics and morphology, i.e., the morphological structure of the spatial regions, which are predicates in languages.
- [11] Indeed, this claim holds even if we attend only to reasoning tasks that only involve the tools of ordinary logic. For example, we seem to prefer conjunction to disjunction. Suppose there are some objects which are either blue or red, and are either square or triangular, and suppose you are asked to place them in two piles, with the ones that are either blue or square in pile A, and the rest in pile B. Although the tasks are equivalent, subjects perform faster when asked to place the red triangles in pile B, and the rest in pile A.
- [12] Of course, this is not to say that the creature's representation individuates all and only the actual objects correctly; its mind does whatever it does, and may be susceptible to biases and errors, like humans are.
- [13] In fact, the result is analogous to using the picturesque image of the model in logical software like Tarski's World as the complete (sentential) description of the set-theoretic model (Barker-Plummer, Barwise, & Etchemendy 2007).
- [14] Camp acknowledges the importance of how representations are processed (2007, pp. 152 & 170), so it's puzzling how, in the absence of any processing theory, any claims about efficiency or tractability are justified.
- [15] For example, it could be that compressed representations are stored in some format like (8), but are of little or no use in as such. Instead, it could be that most of the creature's mental processing involves the individual formulae (3). Intuitively, it would be as if the creature's mind's eye beheld the tightly bundled package (or the box it came in), while unbeknownst to the creature, the busy workers of its processing system were behind it, plucking out individual pieces of the package and reorganizing them in the actual process of reasoning.
- [16] For a very clever and helpful example of how patterns of breakdown can mislead in this respect, see Davies (1987, pp. 450–454).
- [17] I am grateful to an anonymous reviewer for suggesting this reaction.

References

- Anderson, J. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, 85(4), 249–277.
- Armstrong, D. (1973). *Belief, truth, and knowledge*. Cambridge: Cambridge University Press.
- Bach, K. (1987). *Thought and reference*. Oxford: Clarendon.
- Barker-Plummer, D., Barwise, J., & Etchemendy, J. (2007). *Tarski's world*. Stanford, CA: CSLI.
- Bateson, P. (1991). Are there principles of behavioral development? In P. Bateson (Ed.), *The development and integration of behaviour: Essays in honour of Robert Hinde* (pp. 19–39). Cambridge: Cambridge University Press.
- Berwick, R. C., & Chomsky, N. (2011). The biolinguistic program: The current state of its development. In A. M. Di Sciullo & C. Boeckx (Eds.), *The biolinguistic enterprise: New perspectives on the evolution and nature of the human language faculty* (pp. 19–41). Oxford: Oxford University Press.
- Blumson, B. (2012). Mental maps. *Philosophy and Phenomenological Research*, 85(2), 413–434.
- Bogen, J., & Woodward, J. (1992). Observations, theories, and the evolution of the human spirit. *Philosophy of Science*, 59(4), 590–611.
- Braddon-Mitchell, D., & Jackson, F. (1996). *Philosophy of mind and cognition*. Oxford: Blackwell.
- Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, 21, 145–182.
- Carey, S. (2009). *The origin of concepts*. Oxford: Oxford University Press.
- Carruthers, P. (2006). *The architecture of the mind*. Oxford: Oxford University Press.
- Cela-Conde, C. J., & Marty, G. (1998). Noam Chomsky's minimalist program and the philosophy of mind: An interview. *Syntax*, 1(1), 19–36.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. New York: Springer.
- Chomsky, N. (1986). *Knowledge of language*. Westport, CT: Praeger.
- Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
- Chomsky, N. (2002). *On nature and language*. Cambridge: Cambridge University Press.
- Chomsky, N., & McGilvray, J. (Eds.). (2012). *The science of language: Interviews with James McGilvray*. Cambridge: Cambridge University Press.
- Chomsky, N., & Stemmer, B. (1999). An on-line interview with noam chomsky: On the nature of pragmatics and related issues. *Brain and Language*, 68(3), 393–401.
- Churchland, P. (2012). *Plato's camera: How the physical brain captures a landscape of abstract universals*. Cambridge, MA: MIT Press.
- Collins, J. (2007). Meta-Scientific eliminativism: A reconsideration of Chomsky's review of Skinner's *Verbal behavior*. *British Journal for the Philosophy of Science*, 58, 625–658.
- Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge, MA: MIT Press.
- Davies, M. (1987). Tacit knowledge and semantic theory: Can a five percent difference matter? *Mind*, 96(384), 441–462.
- Devitt, M. (2006). *Ignorance of language* (2nd edition). Oxford: Oxford University Press.
- Ebbinghaus, H., & Flum, J. (1999). *Finite model theory*. Heidelberg: Springer.
- Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1997). *Rethinking innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press.
- Fodor, J. A. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.
- Fodor, J. A. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.
- Fodor, J. A. (Ed.). (1998). *Concepts*. Oxford: Oxford University Press.
- Fodor, J. A. (2001). Language, thought, and compositionality. *Linguistics and Philosophy*, 16(1), 1–15.
- Fodor, J. A. (2008). *LOT2: The language of thought revisited*. Oxford: Clarendon.
- Fodor, J. A., & Pylyshyn, Z. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28, 3–71.

- Gallistel, C., & Gelman, R. (2000). Nonverbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences*, 4(2), 59–65.
- Gattis, M. (Ed.). (2001). *Spatial schemas and abstract thought*. Cambridge, MA: MIT Press.
- Goodman, N. (1968). *Languages of art*. Indianapolis, IN: Bobbs-Merrill.
- Griffiths, P. (2009). The distinction between innate and acquired characteristics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Fall 2009 ed.). Retrieved from <http://plato.stanford.edu/archives/fall2009/entries/innateacquired/>
- Hale, K., & Keyser, S. J. (1986). *Some transitivity alternations in English*. Cambridge, MA: Lexicon Project, Center for Cognitive Science, MIT.
- Hale, K., & Keyser, S. J. (1987). *A view from the middle*. Cambridge, MA: Lexicon Project, Center for Cognitive Science, MIT.
- Hatfield, G. (2009). *Perception and cognition: Essays in the philosophy of psychology*. Oxford: Oxford University Press.
- Hodges, W. (1993). *Model theory*. Cambridge: Cambridge University Press.
- Hume, D. (1978). *A treatise of human nature*. Oxford: Oxford University Press.
- Jackendoff, R. (2002). *Foundations of language*. Cambridge, MA: MIT Press.
- Johnson, K. (2001). *Knowledge of word-meaning* (Doctoral dissertation). Rutgers University.
- Johnson, K. (2004). On the systematicity of language and thought. *Journal of Philosophy*, 101(3), 111–139.
- Kosslyn, S. M. (1980). *Image and mind*. Cambridge, MA: Harvard University Press.
- Kosslyn, S. M. (1994). *Image and brain: The resolution of the imagery debate*. Cambridge, MA: MIT Press.
- Laurence, S., & Margolis, E. (2005). Number and natural language. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Structure and contents, volume 1: Evolution and cognition* (pp. 216–235). Oxford: Oxford University Press.
- Marr, D. (1982). *Vision*. San Francisco, CA: Freeman.
- Miller, G. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Nersessian, N. J. (2004). The cognitive basis of model-based reasoning in science. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The cognitive basis of science* (pp. 133–153). Cambridge: Cambridge University Press.
- Peacocke, C. (1986). Analogue content. *Proceedings of the Aristotelian Society, Suppl. Vol. 60*, 1–17.
- Pylyshyn, Z. (2002). Mental imagery: In search of a theory. *Behavioral and Brain Sciences*, 25, 157–182.
- Pylyshyn, Z. (2003). *Seeing and visualizing: It's not what you think*. Cambridge, MA: MIT Press.
- Rescorla, M. (2009). Predication and cartographic representation. *Synthese*, 169, 175–200.
- Robinson, A. (1961). Non-Standard analysis. *Proceedings of the Royal Academy of Amsterdam*, 64, 432–440.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972), 701–703.
- Sober, E. (1976). Mental representations. *Synthese*, 33(1), 101–148.
- Suppes, P. (1977). Is visual space Euclidean? *Synthese*, 35(4), 397–421.
- Thagard, P. (2012). *The cognitive science of science: Explanation, discovery, and conceptual change*. Cambridge, MA: MIT Press.
- Thomas, N. J. (2011). Mental imagery. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Winter 2011 ed.). Retrieved from <http://plato.stanford.edu/archives/win2011/entries/mental-imagery/>
- Tye, M. (2000). *The imagery debate*. Cambridge, MA: MIT Press.
- van Gelder, T. (1990). Compositionality: A connectionist variation on a classical theme. *Cognitive Science*, 14, 355–384.

- van Gelder, T. (1995). What might cognition be, if not computation? *Journal of Philosophy*, 91, 345–381.
- Waskan, J. (2003). Intrinsic cognitive models. *Cognitive Science*, 27, 259–283.
- Waskan, J. (2006). *Models and cognition*. Cambridge, MA: MIT Press.
- Westerhoff, J. (2005). Logical relations between pictures. *Journal of Philosophy*, 102(12), 603–623.
- Xu, F. (1997). From Lot's wife to a pillar of salt: Evidence that physical object is a sortal concept. *Mind & Language*, 12(3/4), 365–392.